ATI X800 XT Preview (page 20) ■ **PC Modder: Tale Of A Mac Modder** (page 33)

A Sandhills Publication

CPU®
COMPUTER POWER USER

See This System On Page 40

Cooligy On Watercooling (page 44) SnapStream Beyond TV 3 (page 76) Big Storage Options (page 16)

“We pointed the industry in the right direction.”

nVIDIA

CEO Jen-Hsun Huang
Dishes On Graphics & Computing
(page 54)

June 2004 • Vol. 4 Iss. 6
$5.99 U.S. $7.99 Canada

CAN-SPAM's Bid For Clutter-Free Inboxes
(page 82)

E-Voting: Can It Ever Work?
(page 48)

(page 64)

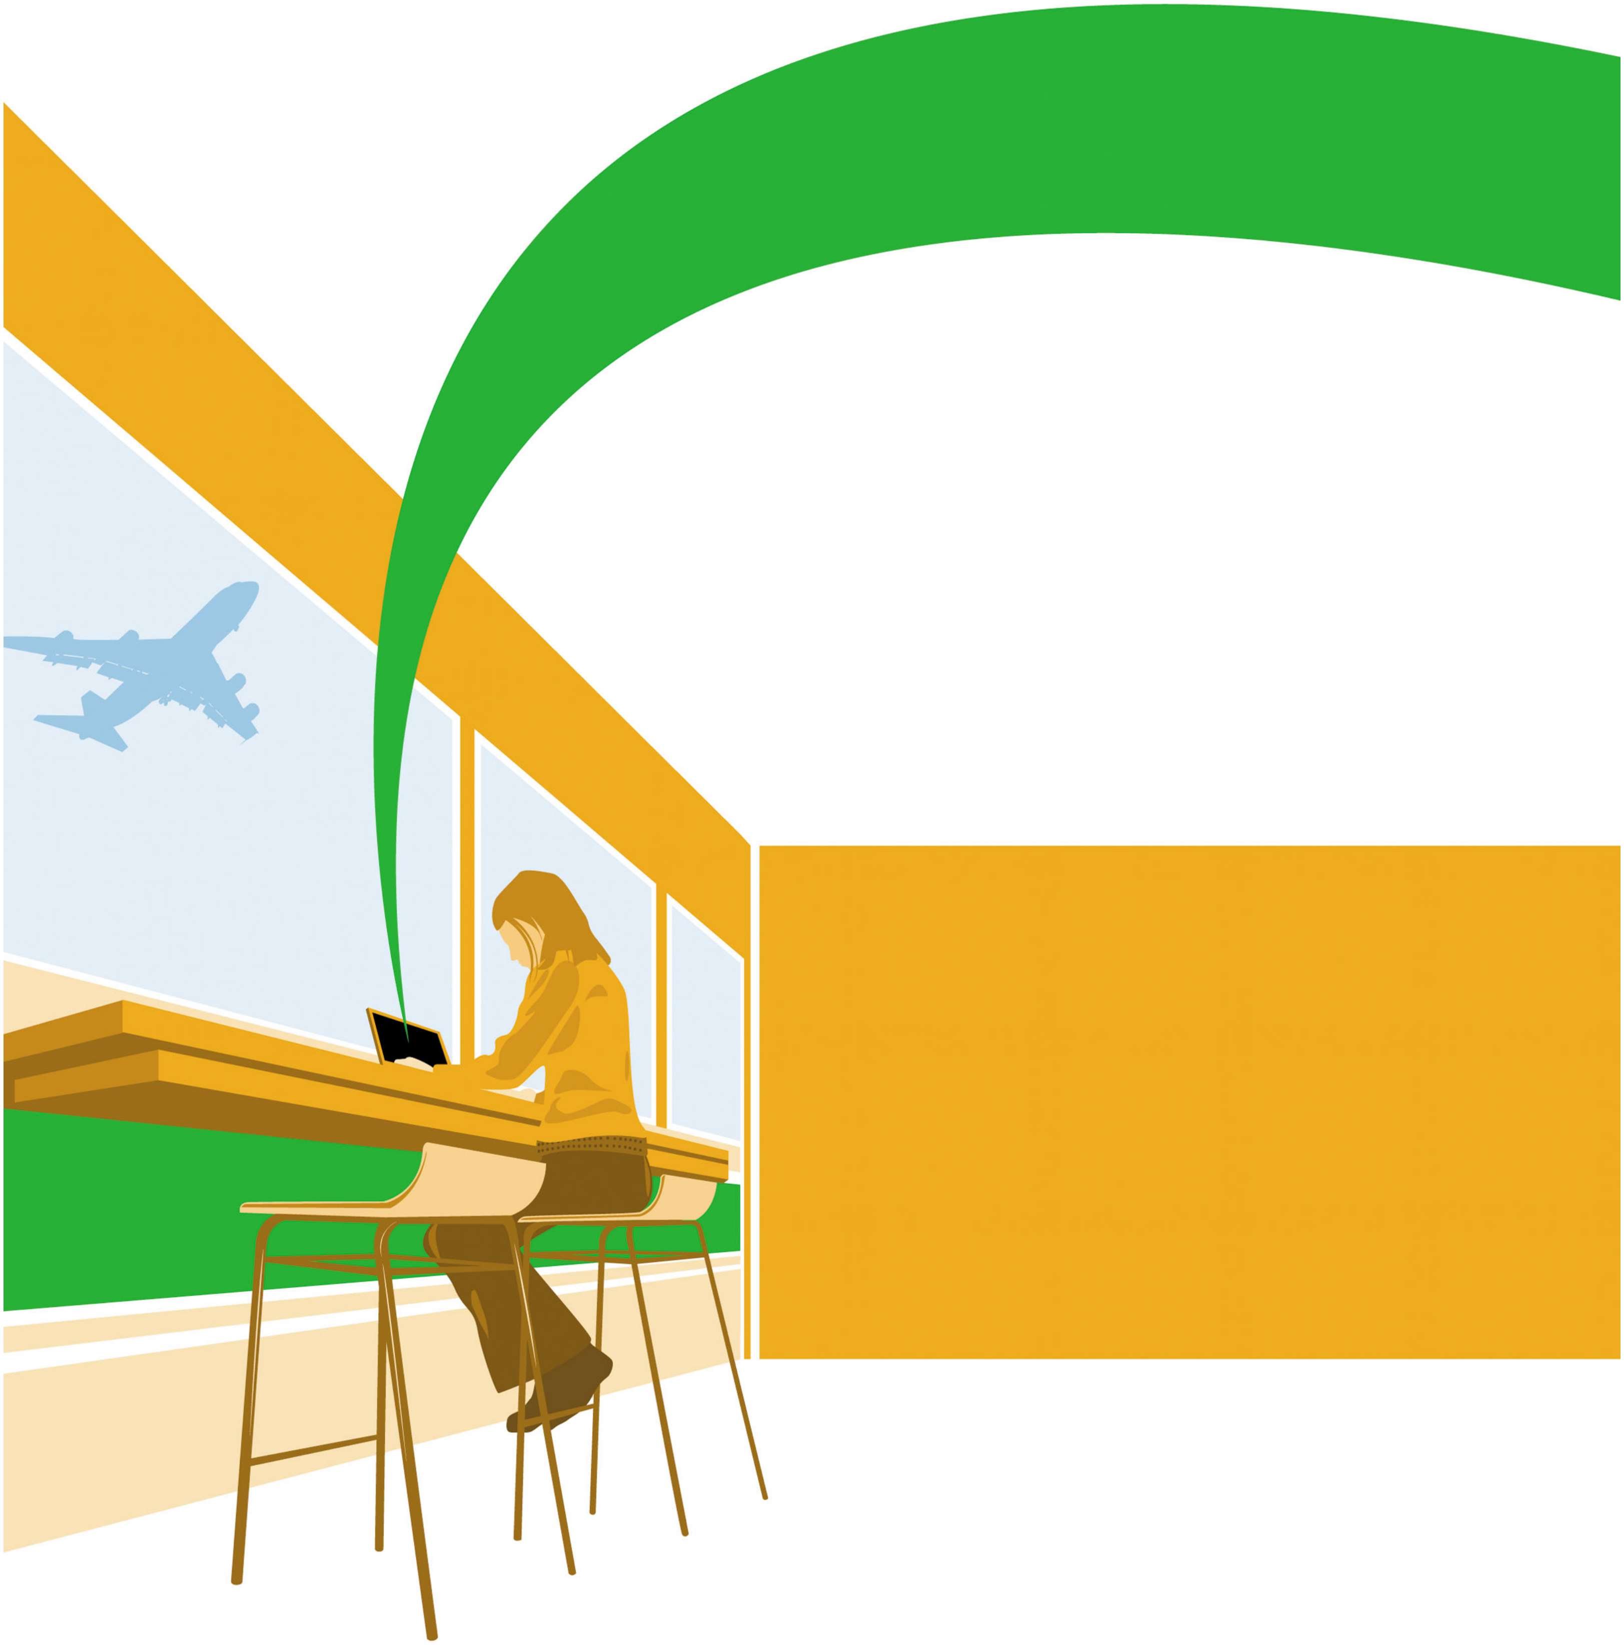

GoToMyPC®

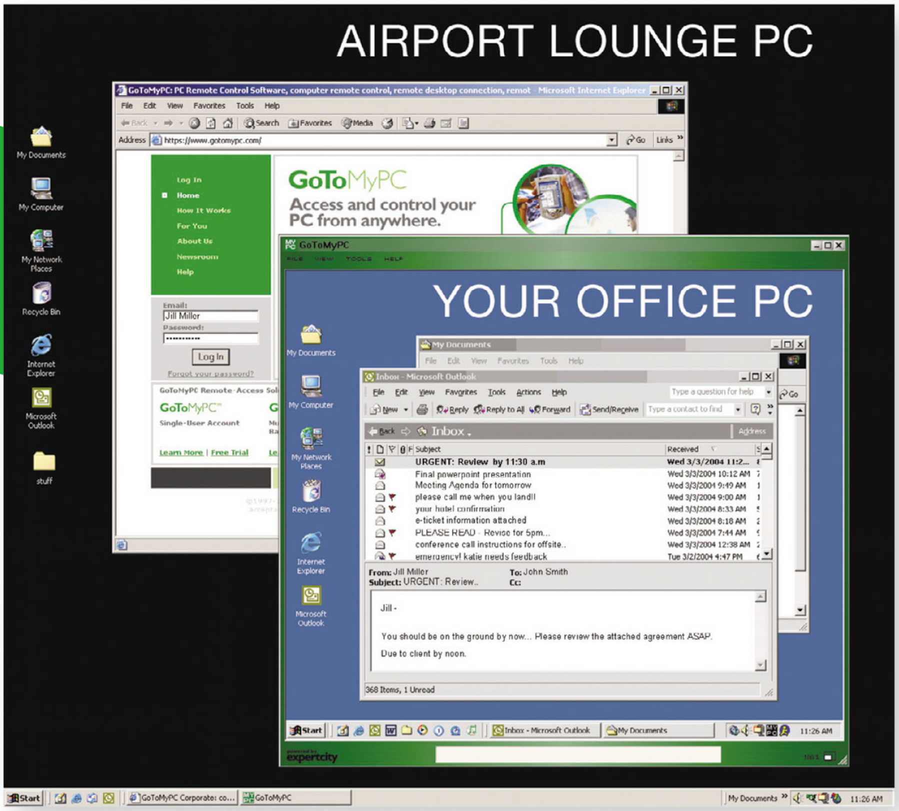

Email: Jill Miller

Subject	Received
URGENT: Review by 11:30 a.m	**Wed 3/3/2004 11:2...**
Final powerpoint presentation	Wed 3/3/2004 10:12 AM
Meeting Agenda for tomorrow	Wed 3/3/2004 9:49 AM
please call me when you land!!	Wed 3/3/2004 9:00 AM
your hotel confirmation	Wed 3/3/2004 8:33 AM
e-ticket information attached	Wed 3/3/2004 8:18 AM
PLEASE READ - Revise for 5pm...	Wed 3/3/2004 7:44 AM
conference call instructions for offsite..	Wed 3/3/2004 12:38 AM
emergency! katie needs feedback	Tue 3/2/2004 4:47 PM

From: Jill Miller **To:** John Smith
Subject: URGENT: Review.. **Cc:**

Jill -

You should be on the ground by now... Please review the attached agreement ASAP.

Due to client by noon.

368 Items, 1 Unread

You're sitting in the **airport lounge**...
your flight is delayed.

You remember there's a presentation
you need to edit on your **office PC**.

You find a comfortable spot in front of a PC in the airport lounge and log on to **gotomypc.com**.

Your office PC desktop instantly appears on the screen in front of you, just as if you were sitting at your desk back at work. With a relaxed smile, you access and work on the presentation you needed and respond to some urgent emails.

With GoToMyPC, using any computer connected to the Internet, or a handheld wireless device, you can always have complete and secure access to everything on your PC from anywhere in the world.

"You'd be nuts not to sign up!" — CNET

Free 30-Day Trial
gotomypc.com/cpu
866.270.1212

CPU
COMPUTER POWER USER

June 2004 • Vol 4 Issue 6

A Look At The Future Of Graphics & Computing

Spotlight

54 **NVIDIA Speaks**
The GPU Giant Comes Clean On What Went Wrong, How They Fixed It, & Where They're Taking Us Next

64 **Weigh-in At The Scales**
NVIDIA's NV40 Is Here & The Battle Is On

CPU
Battle Of The $500 Video Cards
AMD

Frontside

6 **What's Happening**
12 **Digital Economy**
14 **The Saint**
Spyware 2.0

The Experts

Alex St. John
The Saint
page 14

Alex "Sharky" Ross
The Shark Tank
page 29

Joan Wood
Forward Slash
page 89

Anand Lal Shimpi
Anand's Corner
page 28

Rob "CmdrTaco" Malda
The Department Of Stuff
page 88

Kyle Bennett
[H]ard Talk
page 30

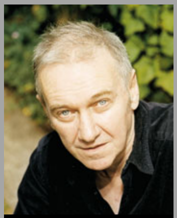

Mike Magee
Shavings From The Rumour Mill
page 102

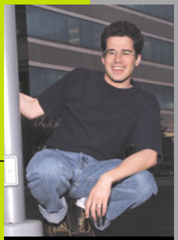

Chris Pirillo
Dialogue Box
page 80

Pete Loshin
Open Sauce
page 81

Copyright 2004 by Sandhills Publishing Company. *Computer Power User* is a trademark of Sandhills Publishing Company. All rights reserved. Reproduction of material appearing in *Computer Power User* is strictly prohibited without written permission. Printed in the U.S.A. GST # 123482788RT0001 (ISSN 1536-7568) *CPU Computer Power User* USPS 020-801 is published monthly for $29 per year by Sandhills Publishing Company, 131 West Grand Drive, P.O. Box 82667, Lincoln, NE 68501. Subscriber Services: (800) 424-7900. Periodicals postage paid at Lincoln, NE. POSTMASTER: Send address changes to *Computer Power User*, P.O. Box 82667, Lincoln, NE 68501.

Did you find the hidden CPU logo on our cover? Turn the page for the answer.

Heavy Gear

15 **Dream Hardware**

16 **Extreme External Storage**
Stockpiling From The Outside In

20 ATI Radeon X800 XT (R420) Preview

21 AMD Athlon 64 FX-53 CPU

22 Intel Pentium 4E 3.4GHz CPU

23 Jetway MiniQ 860TWIN SFF PC

24 MSI Mega 180 SFF Case
Sony Cyber-shot DSC-F828

26 TDK Indi DVD 880N
Epson Stylus Photo R200

27 **Heatsink Hardball**
Aerocool vs. Gigabyte vs. Zalman

28 **Anand's Corner**
GeForce 6–Finally

29 **The Shark Tank**
My Battery Life

30 **[H]ard Talk**
Benchmarks Tell You What?

Hard Hat Area

32 **Mods & Ends**
33 **The ArtMac**
The Tale Of A Mac Modder

40 **Mad Reader Mod**
The Fragmatatron:
A Rusted-Out Beauty

41 Advanced Q&A Corner
44 **X-ray Vision:** Cooligy
Making Watercooling Small & Efficient
48 **White Paper:** E-voting
Soapbox Salvation Or Political Suicide?

Loading Zone

70 **The Bleeding Edge Of Software**
Inside The World Of Betas
72 **Up To Speed**
Upgrades That'll Keep You Humming Along
73 **Enlisting A Video Encoder**
For When It's Time To Compress & Decompress
76 SnapStream Beyond TV3
77 **Apple iLife '04**
Introducing GarageBand With New Versions Of iTunes, iPhoto & More
80 **Dialogue Box**
Hijacked By Spyware
81 **Open Sauce**
Freaky Wiki

Caught In The Web

82 **Can CAN-SPAM Really Can Spam?**
New Laws Are No Match For Unwanted Email
86 **Coder's Corner: XML**
Schema Languages For XML
88 **The Department Of Stuff**
uribl.txt
89 **Forward Slash**
Play Phone

Digital Living

90 **Road Warrior**
The Future Of Wireless, Sendo vs. Microsoft, Unwiring At The Park & More From The Mobile Front
92 **At Your Leisure**
Plug In, Sit Back & Fire Away

Tips & Tricks

97 **Software Tips & Projects**
Slipstream Your Next Windows XP Installation
99 **Warm Up To Penguins**
Sorting Out Regular Expressions

What's Cooking

102 **Shavings From The Rumour Mill**
Chip Speeds Become Irrelevant
103 **Technically Speaking**
An Interview With John Stroozas, Director Of Engineering For Crucial
106 **Under Development**
A Peek At What's Brewing In The Laboratory

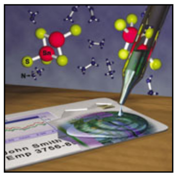

Back Door

108 **Q&A With Jon S. von Tetzchner**
CEO Of Opera Software

Infinite Loops

Strange stats and other oddball items from computing's periphery

84, 95, 98, 100

GREETINGS FROM SAMITLAND

We started laying the groundwork for this month's cover piece a week into the new year. Our focus was bit of a departure from your "typical" issue of *CPU*, and we needed to plan carefully to ensure we captured something valuable for you, our reader. I mulled many CEOs when considering the focus and kept coming back to NVIDIA and CEO/co-founder Jen-Hsun Huang. Thinking back to 1995, I still recall being excited about the NV1, which was to be shipped on a card called the Diamond Edge, which I eventually received for testing. I was impressed with its feature-set and the 3D-accelerated Virtua Fighter bundled with the card. However, in a massively bad strategic move, the Edge was not DirectX-friendly (granted, DX was new on the scene) and was gone before you could say, "Boo!" Bye-bye. Au Revoir. Arrivederci. Adios. In fact, that failure actually makes the NV30 look like a prizewinner. But for every down, the company has managed to come back stronger than ever. You see that clearly with the GeForce 6 series and in other areas of the company. If that trend continues the way it has (and I don't see why it shouldn't), I'm intrigued to see where the company is a decade or more from now. My personal feeling is that NVIDIA has yet to morph into a butterfly and that the best is yet to come. Being a larva in the cocoon is an exciting place to be.

So without further ado I'd like to introduce you to NVIDIA's main man, Jen-Hsun Haung. We've also got the VP Of Hardware Engineering and co-founder, Chris Malachowsky, on the line. Over the past couple of months we have spent many hours on the NVIDIA campus in Santa Clara harassing these kind folks, and they have been very forthcoming, so kudos to both. We appreciate your efforts. Together, both gentlemen will give you some powerful insight into not just graphics but computing and where things are headed. You can read what they have to say on page 54. As if that's not enough, there's a ton of information to get you up to speed with NVIDIA's latest GPU: NV40 (GeForce 6800). Yeah, it's much cooler than the NV1. Seriously, after the madness that was NV3x, it's good to see the company back on track with NV40. Early indications are that it's a home run.

Of course, we have tons of other material for your reading pleasure. With all our NVIDIA-focused coverage this month, it was nice to have ATI pitch in with a little product of its own. It's called the ATI X800 XT (page 20) and we've got the scoop. I've spent a couple of months playing with SnapStream's Beyond TV 3 (page 76), and I love it. I'm not here to endorse products, but wow, it's seriously cool. C'mon, the ability to blow by advertising? Fantastic! Another very cool package you'll find on the following page is Apple's iLife '04 (with GarageBand). Speaking of Apple, be sure to check out the G4 upgrade/mod piece on page 33. Last but not least are the upcoming elections. E-voting is going to be increasingly important as we go forward; those of you ready to rock the e-vote may wish to check out page 82. There's no time like now to get on the ball with technology that will affect a country's leadership, right? OK, but I'm going to get off my soapbox before I even get on it.

Remember, we want to put together a better product for you. It's really that simple. So let us know what you like and don't like about this issue. Now, I've got to get home and tidy before Julie drives in from Colorado. Until next month!

Samit G. Choudhuri, Publication Editor, *CPU*

Sandhills Publishing

Web Services
(For questions about our Web site.)
webhelp@cpumag.com
(800) 368-8304

Customer Service
customer-service@cpumag.com
Computer Power User
P.O. Box 82667
Lincoln, NE 68501-2667

Hours
Mon. - Fri.: 8 a.m. to 8 p.m. (CST)
Sat.: 8 a.m. to 4 p.m. (CST)

To Place An Order Or Change An Address
(800) 733-3809
FAX: (402) 479-2193

Subscription Renewals
(800) 382-4552
FAX: (402) 479-2193
http://www.cpumag.com

Authorization For Reprints
(800) 334-7458
FAX: (402) 479-2104

Editorial Staff
editor@cpumag.com
FAX: (402) 479-2104
131 W. Grand Drive
Lincoln, NE 68521

Advertising Staff
(800) 848-1478
120 W. Harvest Dr.
Lincoln, NE 68521
FAX: (402) 479-2193

The Gang

Editorial Staff: Ronald D. Kobler / Samit Gupta Choudhuri / Jennie Schlueter / Blaine Flamig / Corey Russman / Rod Scher / Christopher Trumble / Calvin Clinchard / Kimberly Fitzke / Katie Dolan / Raejean Brooks / Rebecca Christensen / Tara Weber / Sally Curran / Michael Sweet / Katie Crooks / Shawn Kupfer / Dana Montey / Nate Hoppe / Sheila Allen / Heather Svoboda / Jennifer Hansen / Trista Kunce / Marty Sems / Chad Denton / Nathan Chandler / Kylee Dickey / Josh Gulick / Andrew Leibman / Eric Reed **Web staff:** Missy Fletcher / Mary Lafferty / Dorene Krausnick / Nick Ray **Customer Service:** Travis Brock / Alisha Lamb / Brandie Humphrey / Becky Rezabek / Lana Matic **Subscription Renewals:** Liz Kohout / Connie Beatty / Matt Bolling / Patrick Kean / Charmaine Vondra / Miden Ebert / Kathy DeCoito / Stephanie Contreras / Nicole Buckendahl **Art & Design:** Lesa Call / Fred Schneider / Carrie Benes / Ginger Riley / Sonja Warner / Leigh Trompke / Aaron Weston / Aaron Clark / Lori Garris / Jason Codr / Andria Schultz / Erin Rodriguez / Lindsay Anker **Newsstand:** Garth Lienemann / Kelly Richardson / Chris McGreer / Jeff Schnittker **Advertising Sales:** Grant Ossenkop / Cindy Pieper / Brooke Wolzen / Eric Cobb **Marketing:** Mark Peery / Marcy Gunn / Heather Mentgen / Amber Coffin

Gotcha. Here it is.

AMD

AMD 64 Athlon FX

Introducing the AMD Athlon™ 64 FX processor. Take your system to extremes.
Double the data path from 32- to 64-bit and you more than double the thrill factor. Uninterrupted, ear-splitting, streaming audio and rich, razor sharp video make your pad a launching pad. What's more, you get all the power you need to edit, mix, and model your own digital creations with memory to spare. Prepare to blow minds. Get a dose of the AMD Athlon 64 FX edge at www.amd.com/amdathlon64fx

© 2003 Advanced Micro Devices, Inc. All rights reserved. AMD, the AMD Arrow logo, AMD Athlon, and combinations thereof are trademarks of Advanced Micro Devices, Inc.

Compiled by
Steve Smith

Microsoft's Eye For The Tech Guy

Microsoft's Angels: Immersion (blue), Night Vision (green), and Groovy (orange).

What's next? Bill Gates in bell bottoms, sideburns, and a perm? After years of unremarkable mouse designs, Microsoft has piled on the makeover bandwagon in an effort to add style and color to our desktops. The new Special Edition Microsoft wireless mice have unique color schemes and textures, including one dubbed Groovy that seems to have been inspired by the Scooby Doo Mystery Machine bus. The Night Vision model appeals to cyberspace culture, while Immersion follows an undersea motif.

Priced between $44.95 and $54.95, all models include the latest optical tracking technology from Microsoft, as well as the recently introduced Tilt Wheel, which adds lateral movement to the scroll wheel. These Special Edition mice are available for a short time, but consumers are invited to cast their vote on the next round of designs issued for the series at www.microsoft.com/hardware/personality, where you can also get Desktop wallpaper to match the designs. ▲

Now That's A Chip!
IBM Demonstrates Its Power5

Imagine trying to pop this one onto an ATX PC motherboard. The upcoming IBM Power5 MCM is aptly dubbed by the company a "multichip monster" (thus MCM). The stats alone are staggering for the package, according to a preview from TheInquirer.net: four Power5 CPUs with a combined cache of 144MB and 276 million transistors in a package that requires about 5,400 pins, with 2,313 dedicated to passing signals and the rest for powering the darn thing.

IBM recently demonstrated the chip for the first time publicly, juggling multiple operating systems in what the company calls "virtual micropartitions." The Power5 has been up and running in IBM's labs for a year already, and it will occupy the company's eServer platform and storage servers when they roll out commercially sometime in 2004. The Department of Energy has already ordered 12,000 of these puppies to run its ASCI Purple project, which will be a supercomputer capable of 100 teraflops of processing power. We don't even want to think about the heatsink needed for this kind of arrangement. ▲

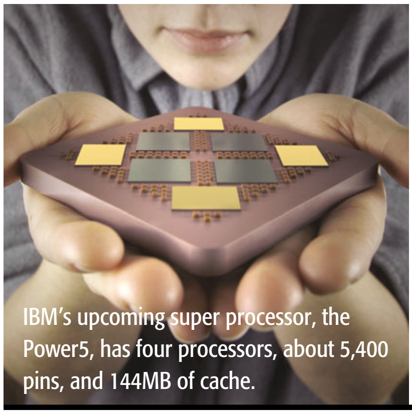

IBM's upcoming super processor, the Power5, has four processors, about 5,400 pins, and 144MB of cache.

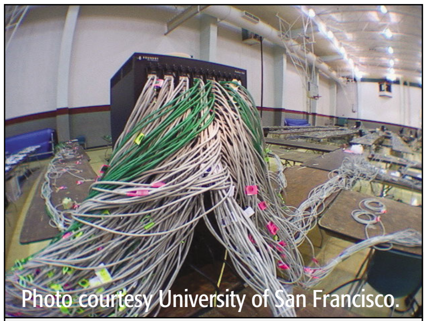

Photo courtesy University of San Francisco.

Flash Mob Computing
Suffers Boot Problems

A noble attempt by researchers and students to generate supercomputing power from 700 networked PCs at the University of San Francisco missed the mark in early April. Members of the "Flash Mob 1" project (www.flashmobcomputing.org) gathered the collection of volunteered PCs and notebooks into a gym and networked them to participate in a massive calculation called the Linpack benchmark. Two portable generators were brought in to power this attempt to create a supercomputer capable of performing 500 billion operations per second, which would qualify it as one of the top 500 supercomputers on record. Some of the PCs proved unstable, however, and in the end, the Flash Mob generated only 180 billion operations per second. Organizers were undaunted, insisting that the experiment demonstrated the potential of democratic computing, joining PCs volunteered by citizens into massive computer power. ▲

Photo courtesy University of San Francisco.

It's Official:
World's Smallest Hard Drive

Joining the ranks of the farthest eyeball popper (that would be Kim Goodman of Chicago) and the most condolences sent on the Internet (that would be 580,000 to the British Royal site upon the death of Princess Diana), Toshiba's 0.85-inch hard drive was recognized by the Guinness World Records Book as the world's smallest hard drive. The drive will be available in late 2004 initially in 2GB to 4GB capacities. It measures a quarter of the size of Toshiba's current reigning midget, the 1.8-inch drive. The single platter design will deliver a rotational speed of 3,200rpm and weigh less than 10 grams. Toshiba anticipates the device will be used in bringing extended storage capabilities to cell phones, PDAs, and digital cameras. ▲

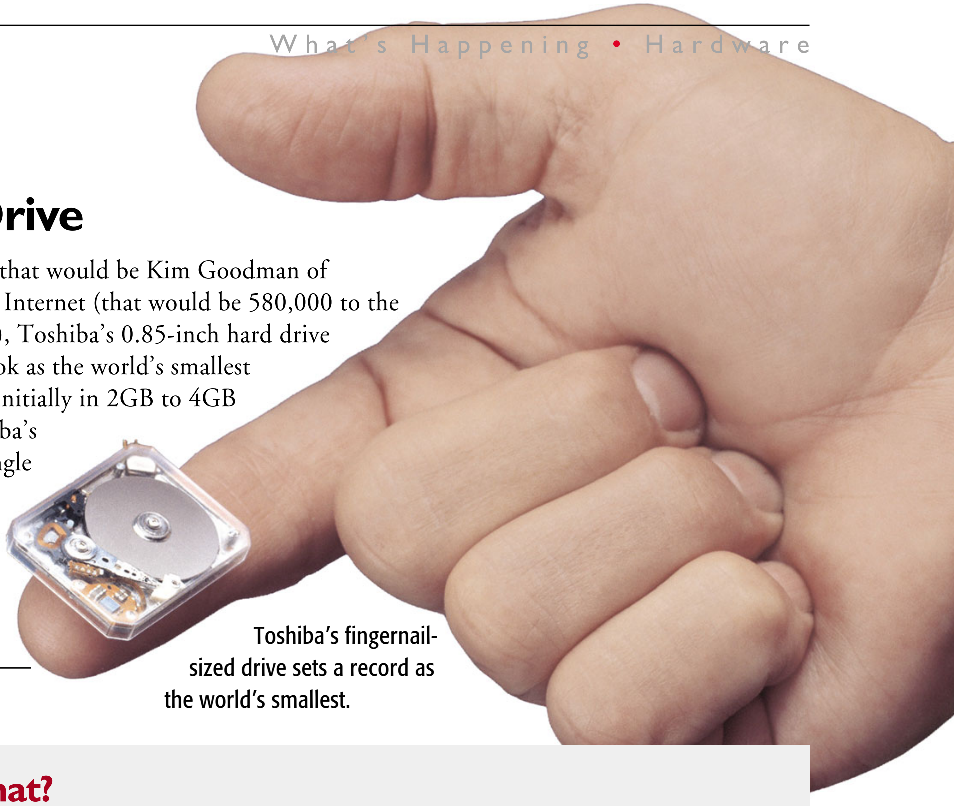

Toshiba's fingernail-sized drive sets a record as the world's smallest.

Hardware Mole

Do GPS Slacks Come With That?

Being a gadget nerd doesn't mean you have to look like one. SCOTTeVEST (www.scottevest.com) just introduced the TEC Sport Jacket, a $250 jacket for the gadget obsessed. It boasts 14 hidden pockets designed specifically to hold PDAs, cell phones, and even earbuds. The PAN (Personal Area Network) is a series of cloth channels that direct any connective wiring among your devices. Collar loops hold headphones/earbuds in place. Have no fear, however, as all of this techie overkill is well-hidden by worsted wool that won't sag under the weight of all of this hardware and won't overheat even with every lithium-ion battery in your gadget collection cooking at full blast. ▲

Blu-ray For PS3?

Speculating on the capabilities and feature sets of the next generation of console game hardware has become an amateur sport in the game press as Sony, Microsoft, and Nintendo keep super-secret their plans and everyone hangs on each morsel of suggestive information. In an interview with a Japanese PC magazine, Sony's manager of Blu-ray disc R&D let slip that the super-dense storage format may end up in the PlayStation 3, expected in 2006. With 25GB to 50GB capacity, Blu-ray discs (BD-ROM) would give game developers unprecedented, massive amounts of space for game data or (yikes) longer cut scenes than any of us could bear. ▲

Kill Faster, Jump Higher

Strange but true . . . and inevitable: Gamers now have a sports glove designed specifically for the stress and strain of extended console controller use. The two-fingered DigiGlove ($13.95 from IntroMark at 800/851-6030) has extra padding at the forefinger and thumb to prevent soreness and calluses. The entire design allows for free airflow because no one wants a sweaty trigger finger, and there is optimal flexibility at the forefinger knuckle so the glove doesn't interfere with your fragging speed. The inventor boasts that the glove also "reduces the risk of White Finger Vibration." We don't know what that is, exactly, but we imagine it's the gamers' version of athlete's foot. ▲

The DigiGlove is the latest in game addict gear.

Compiled by Dean Takahashi

Intel Enters Race To Put 3D Graphics On Handhelds

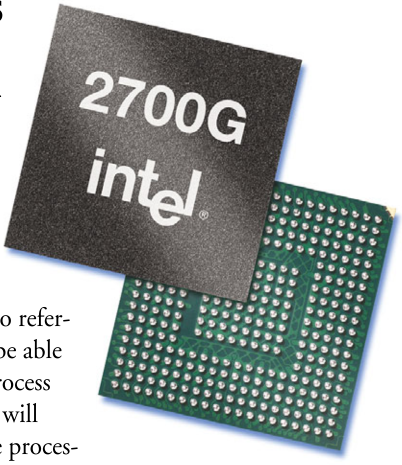

Intel isn't a player in standalone graphics chips for computers, but the world's biggest chipmaker hopes its Marathon multimedia accelerator will snare it a spot in handhelds and cellular phones. Robert McNair, director of marketing for Intel's handheld products, says the Marathon chip can process 944,000 triangles per second with low power consumption, making it ideal for portable applications. The chip will be part of the Carbonado reference design for a personal digital assistant and will be able to drive a 3- or 4-inch display. The chip will also process MPEG-2 video at 30 frames per second. Marathon will be a co-processor for a handheld using the Bulverde processor, which is the next-generation chip in the Xscale handheld microprocessor family. Intel designed Marathon from the ground up, McNair says. "We see that what happened 10 years ago in PCs is starting to happen now in handhelds, and we want to be there," he says. "What differentiates us is we can make the whole platform." McNair says the chips will be shipping soon and that devices using them will debut later this year. ▲

XGI Hopes To Take On NVIDIA & ATI In 3D Graphics

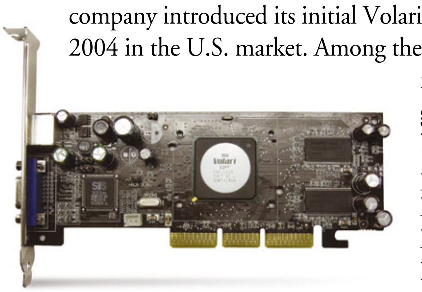

NVIDIA and ATI have already dispatched dozens of rivals who have tried to take them on in the 3D graphics market, but the startups just keep coming. Taiwan-based XGI Technology is the latest to enter the 3D market with its family of Volari 3D chips. The company introduced its initial Volari chips in the fall of 2003 in Asia and in the spring of 2004 in the U.S. market. Among the Volari's unique features is the ability to put two of its Volari V8 Ultra chips in a single graphics card to get higher performance. Wei-wei Yin, director of business development at XGI, said the company plans to have a full product line, spanning from its low-cost Volari V3XT chip coming out this fall to a high-end chip that will be competitive with ATI's R420 chip. Third-party analysts haven't yet weighed in on the performance of XGI's chips, but the company is hoping that OEMs will like the competition enough to give XGI a shot. "We also think fans will like a third vendor coming to the market," says Yin. ▲

IBM Brings Open Source To Its PowerPC Chips

IBM gave open-source software a jumpstart with an endorsement of Linux, and now it hopes to do something similar with its microprocessors. The company says it will allow other companies to license its PowerPC architecture so that they can build custom chips around the PowerPC designs. To date IBM has designed and built most of its PowerPC chips. But under the new approach, others can license PowerPC cores and include them in their own chips. The business model resembles that of ARM Ltd. and MIPS Technologies, who will be prime competitors. The first licensee will be Sony, which is expected to use PowerPC cores in a variety of electronics products. Sony has already signed up with IBM and Toshiba to create the Cell microprocessors that are expected to be used in the PlayStation 3. ▲

Watching The Chips Fall

Here is pricing information for AMD and Intel CPUs.

Released	Original Price	Current Price (04/06/04)	Last Month's Price
AMD Athlon XP 2600+			
8/21/2002	$297**	$79	$81
AMD Athlon XP 2700+ 333MHz FSB			
11/14/2002	$349**	$97	$95
AMD Athlon XP 2800+ 333MHz FSB			
11/14/2002	$397**	$106	$112
AMD Athlon XP 3000+ 333MHz FSB			
2/10/2003	$634*	$143	$152
AMD Athlon XP 3000+ 400MHz FSB			
5/26/2003	$280*	$165	$199
AMD Athlon XP 3200+ 400MHz FSB			
5/26/2003	$464**	$205	$206
AMD Athlon 64 2800+			
1/6/2004	$193**	$173	$180
AMD Athlon 64 3000+			
9/23/2003	$218**	$211	$205
AMD Athlon 64 3200+			
9/23/2003	$417**	$270	$270
AMD Athlon 64 3400+			
1/6/2004	$417**	$405	$400
AMD Athlon 64 FX-53			
3/18/2004	$733**	$745	$733**
Intel Pentium 4 2.53GHz 533MHz FSB			
5/6/2002	$637	$168*	$168
Intel Pentium 4 2.6GHz 800MHz FSB			
5/21/2003	$228*	$177*	$177*
Intel Pentium 4 2.66GHz 533MHz FSB			
8/26/2002	$401**	$161*	$161*
Intel Pentium 4 2.8GHz 533MHz FSB			
8/26/2002	$508**	$162*	$170
Intel Pentium 4 2.8GHz 800MHz FSB			
5/21/2003	$278*	$179*	$175
Intel Pentium 4 2.8GHz 800MHz FSB 1MB cache			
2/2/2004	$178**	$184*	$179
Intel Pentium 4 3.06GHz 533MHz FSB			
11/14/2002	$658*	$209	$209
Intel Pentium 4 3.0GHz 800MHz FSB			
4/21/2003	$417**	$208*	$219*
Intel Pentium 4 3.0GHz 800MHz FSB 1MB cache			
2/2/2004	$218**	$233*	$259*
Intel Pentium 4 3.2GHz 800MHz FSB			
6/23/2003	$637**	$280*	$278*
Intel Pentium 4 3.2GHz 800MHz FSB 1MB cache			
2/2/2004	$278**	$278	$278
Intel Pentium 4 3.4GHz 800MHz FSB			
2/2/2004	$417**	$416*	$420*
Intel Pentium 4 Extreme Edition 3.2GHz 2MB cache 800MHz FSB			
11/3/2003	$925**	$949*	$859
Intel Pentium 4 Extreme Edition 3.4GHz 2MB cache 800MHz FSB			
2/2/2004	$999**	$1,079*	$1,079*

* Retail price
** Manufacturer's price per 1,000 units
Other current prices, if indicated, are lowest OEM prices available through Pricegrabber.com.

Sun & Microsoft Kiss & Make Up

Like old friends . . . sort of. Scott McNealy, CEO, Sun Microsystems, and Steve Ballmer, CEO, Microsoft, smile for the cameras as they announce settlement of legal disputes and a 10-year cooperative pact.

Microsoft and Sun Microsystems have been the Hatfields and the McCoys of the tech world for as long as anyone can remember: bickering, sniping, and finally going to court over patent and antitrust disputes. Sun had been instrumental in the European Union's recent $613 million judgment against Microsoft for unfairly leveraging its Windows monopoly to promote its server and media player technologies. In a landmark agreement between the two goliaths, Sun receives $1.95 billion in payments from Microsoft to settle past patent infringement suits, to resolve antitrust claims, and as royalties on future use of technologies. The companies will share technologies so that their servers can communicate and interoperate more effectively, and they will even collaborate on Sun Java and Microsoft .NET.

Many feel this deal signals a new spirit of cooperation on Microsoft's part, as it appeals the EU decision and also faces down another antitrust suite from RealNetworks over media players. Steve Ballmer, Microsoft CEO, says of the Sun agreement, "Our companies will continue to compete hard, but this agreement creates a new basis for cooperation that will benefit the customers of both companies." ▲

Copycats, You're So Busted

The instant availability of billions of pages of text on the Internet has made plagiarism as easy as cut and paste for all too many students and businesspeople. The iThenticate system from iParadigms (www.ithenticate.com) turns the tables on copycats by offering a Web service that compares a submitted document to a vast digital database of copyrighted material such as books, journals, and Web sites to detect plagiarized content. The system also lets content owners explore whether and how others are misusing their intellectual property. The U.N. Security Council is already using iThenticate to check the originality of the many reports written in its name. It turns out that plagiarism is good business—for the watchdog. Businesses pay $1,000 annually for the service and $10 for every page submitted and analyzed. ▲

Software Shorts

Here are a few software tidbits we found just before press time.

PeopleSoft's Vaporware Lawsuit

Like Monica Lewinsky and Barney the purple dinosaur the term "vaporware" now seems like a relic of the '90s, when tech companies tried to out-hype one another with promised software and feature sets that never actually appeared on the market. The term was revived in March by a $510 million lawsuit against PeopleSoft by Cleveland State University. PeopleSoft software runs much of the school's internal financial aid, admissions, and student scheduling systems. The university claims that PeopleSoft representatives lied about features in the final product, which produces frequent problems in processing student data. PeopleSoft vigorously denies the allegations, but the university says, ". . . the Financial Aid in particular was at the time little more than vaporware in an extremely early stage of development." ▲

Poor Bill

Well, not really, but for the first time we can recall, the software giant's position as the world's richest man is being challenged. While the venerable *Forbes* annual list still taps Bill Gates as the wealthiest, Swedish business weekly *Veckans Affarer* says that the founder of furniture retailer IKEA, Ingvar Kamprad, wears the crown, with a fortune valued at $53 billion, about $6 billion more than *Forbes* says Gates is worth. Skeptics are dickering over how these fortunes are being calculated (Kamprad ranks 13 in the *Forbes* list), but the Swedish publication argues that the lower value of the dollar on foreign exchanges explains some of the shift. ▲

Vigilante Virus?

In late March, a group of Russian virus writers calling themselves the eMule Project issued the NetSky-Q worm that sent distributed denial of service attacks on large and small P2P home pages such as Kazaa and eDonkey, as well as a host of smaller warez sites. The group claimed to be on the side of the good guys, saying that it hoped to educate users about the ills of hacking and copying protected intellectual property. The attacks were scheduled to occur through the middle of April, and by early in the month, eDonkey's home page had been disrupted, although Kazaa had managed to repel the DDoS attacks. No word on whether anyone anywhere learned a single thing about content rights from this worm. ▲

Prof. Google's Lab Set

If you are still just playing with the Autofill settings on your Google Toolbar, you haven't even scratched the surface of what one of the most innovative search companies on the planet has in store for us in the future. Anyone can get a peek by tapping into Google Labs (labs.google.com), where the company regularly posts novel ways of searching the Web and viewing results. For instance, "Personalized Web Search" lets you set up a profile of your interests so that Google can align the search results on any terms you submit more closely with your personal tastes. Froogle Wireless lets you run product searches through Google's shopping engine from your cell phone. One of the coolest experiments is the Web Viewer, which delivers search results as slideshow images.

Google Labs has a good track record. This is where the company beta tested tools, such as local searching and news alerts, that eventually went official. ▲

A beta program in the Google Lab lets you search for product vendors via Froogle on your cell phone.

Current Events . . . In The Abstract

This just in . . . brown rectangles. And in international affairs . . . green squares. For those who think RSS feeds and cell phone alerts are the cutting edge of news delivery, try clicking over to Newsmap (www.marumushi.com/apps/newsmap/newsmap.cfm) for a truly novel way of viewing current events. Using treemap theories and technologies developed at the University of Maryland, Newsmap translates the categories and quantity of news coursing through the Google News portal into relative sizes and colors to give a graphic representation of the flow of news.

Looking like a cross between a hypertext newspaper and a Mondrian painting, the Newsmap lets you drill into major news topics such as technology, sports, and health and sort the news by 10 major nations. Treemaps can be applied to many types of data, from stock market activity to blogs. According to the UMD Human-Computer Interaction Lab behind the approach, treemaps are "space-constrained visualizations of hierarchical structures."

Those people who really want to understand what that means can download a demo of the treemap software at www.cs.umd.edu/hcil/treemap/TMDemolicense.shtml to see how to translate their own pile of information into pretty shapes and colors. ▲

Newsmap visually displays the aggregated coverage that flows through the Google News portal.

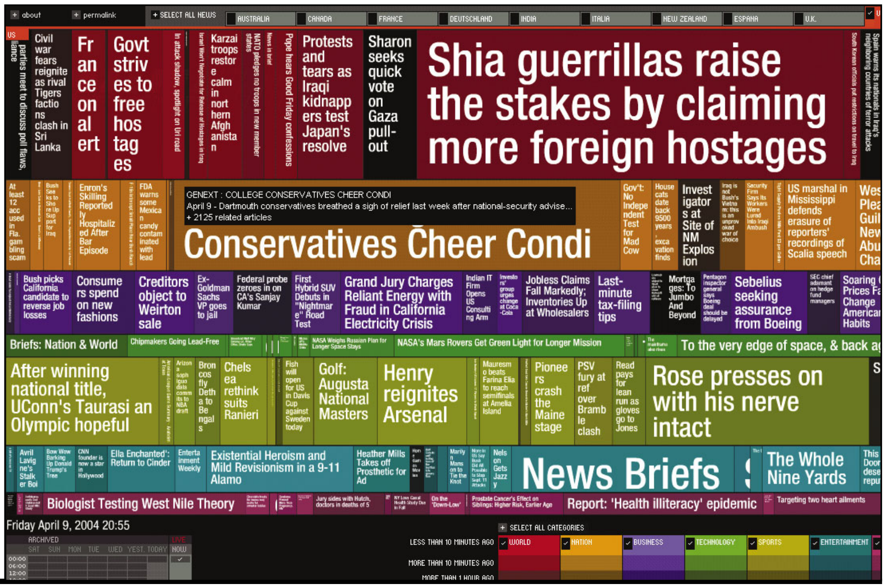
about
permalink
SELECT ALL NEWS
AUSTRALIA
CANADA
FRANCE
DEUTSCHLAND
INDIA
ITALIA
NEW ZEALAND
ESPAÑA
U.K.
Shia guerrillas raise the stakes by claiming more foreign hostages
Conservatives Cheer Condi
Govt strives to free hostages
Protests and tears as Iraqi kidnappers test Japan's resolve
Sharon seeks quick vote on Gaza pull-out
Briefs: Nation & World
Rose presses on with his nerve intact
After winning national title, UConn's Taurasi an Olympic hopeful
Henry reignites Arsenal
Golf: Augusta National Masters
News Briefs
The Whole Nine Yards
Biologist Testing West Nile Theory
Report: 'Health illiteracy' epidemic
To the very edge of space, & back again
Friday April 9, 2004 20:55
ARCHIVED
SELECT ALL CATEGORIES
WORLD
NATION
BUSINESS
TECHNOLOGY
SPORTS
ENTERTAINMENT
LESS THAN 10 MINUTES AGO
MORE THAN 10 MINUTES AGO
MORE THAN 1 HOUR AGO

"You're Wired!"

The Internet Finally Gets "Real"

We've lived through "Survivor," cringed over "Joe Millionaire," and marveled at Donald Trump's hair in "The Apprentice." The Internet finally gets its own reality show in AOL's "The Startup." Co-produced with *Entrepreneur Magazine*, the year-long series follows the ups and downs of diverse new businesses, from a girl's clothing store to a Colorado cavern park. Unlike most TV reality shows, however, the online video vignettes are only the beginning of the drama in "The Startup." Company executives will also host their own online blogs to keep fans apprised of their day-to-day progress, and the AOL cameras will chronicle the nuts and bolts of their business acumen, from budgeting to personnel management. ▲

Broadband Reaches Critical Mass

High-speed Web access has penetrated the American market so deeply and quickly that it is now a mainstream medium that will accelerate the adoption of a host of new technologies, according to In-Stat/MDR. By the end of 2003, nearly 27 million homes and businesses subscribed to broadband, including one in every five households. Senior Analyst Daryl Schoolar says that the subscriber bases for broadband are now large enough to make technologies such as VoIP, Internet links to home entertainment systems, and online gaming explode. "This starts a cycle where growth in both broadband and applications feed the growth of each other," he says.

According to research firm eMarketer, more than half of U.S. households will enjoy high-speed Internet access by 2005. Worldwide broadband markets are expanding even faster, with 98.8 million subscribers in 2003 and 245 million likely by 2007. China already has the third largest broadband subscriber base with 12 million, but as the technology spreads in the world's largest country, it will become the world powerhouse in high-speed Web usage. ▲

Broadband Subscribers Worldwide, 1997, 2000, 2003, 2006, 2007 (In Millions)

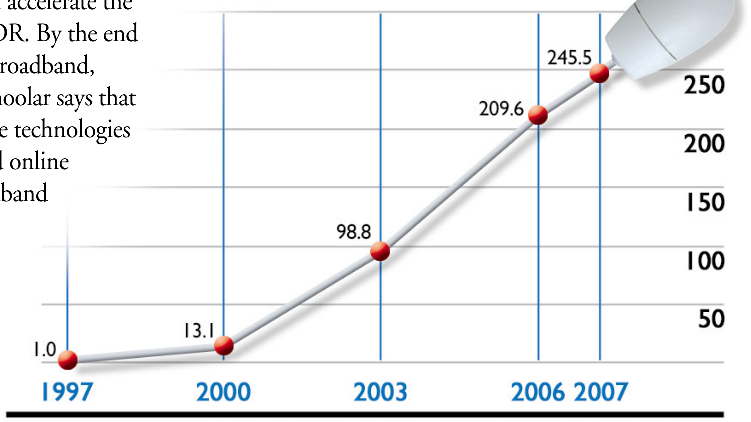

*Note: * Includes all residential and business subscribers but does not include wireless LAN, Bluetooth or 3G mobile braodband subscribers*
Source: eMarketer, March 2004

New On The 'Net

Looking for some new surfing destinations? Here's a sampler of the many sites that recently hit the Web.

Cubicle Rage Gets A Voice

Loathing that jerk in the next cubicle is one of the twisted joys of corporate employment. At iworkwithfools.com, you can share that self-satisfied conviction that you're the only one in the office who knows what she is doing. The site lets you anonymously post complaints and unbelievable tales of workplace stupidity . . . and revenge. Other members commiserate in well-trafficked comments pages. The site's mission couldn't be simpler: "To show the world what a complete bunch of morons we all work with." ▲

Cuban's Blog

Well, that's what Cuban himself says at his new Blog Maverick site (www.blogmaverick.com). The former CEO of Broadcast.com is now the owner of the Dallas Mavericks pro basketball team and very frank narrator of his own online journal. One item, "I Suck," apologizes to fans for the previous night's loss, with Cuban taking the blame because he tried to face down superstition and not wear his lucky shirt to the game. When he isn't raging at officials, Cuban is recounting his involvement in an upcoming ABC reality show, "The Benefactor." The show will chronicle how some lucky person makes use of $1 million, which Cuban will give him or her. ▲

Ken & Barbie Date Online?

We thought that Ken and Barbie had broken up, but apparently they have gotten together on a Web venture that reviews online dating sites. TryInternetDating.com is a consumer guide to one of the Web's hottest trends: digital matchmaking. Hosts Ken and Barbie are old friends (just friends) who discovered that they were making some of the same mistakes when using online dating services. In April they opened a reviews site dedicated to the platform, inviting visitors to share their opinions and embellishing the site with hints and tips for navigating the new world of digital matchmaking. ▲

BIOS Upgrades Available Online

Compiled by Kevin Savetz

Before you send another motherboard to the landfill, consider upgrading the BIOS and giving your PC a new outlook on life. Here are a few recently released upgrades. Subscribers can check out www.cpumag.com/cpujun04/bios to see our entire upgrade list.

Motherboard	Date Available	URL
ABIT NF7-S	03/29/2004	www.abit-usa.com/downloads/bios/bios_revision.php?categories=1&model=6
AOpen MK77M II	04/08/2004	download.aopen.com.tw/downloads/userdownload_List.asp?RecNo=7981&Model=MK77M%20II
AOpen AK79D-400VN	04/08/2004	download.aopen.com.tw/downloads/userdownload_List.asp?RecNo=7982&Model=AK79D-400VN
ASUS SK8N	04/05/2004	www.asus.com/support/download/download.aspx
Chaintech 9VJL5	03/23/2004	www.chaintechusa.com/tw/eng/Download/dl_desc.asp?DCSNo=4&PISNo=195
Gigabyte GA-8S661FXM-F	04/01/2004	tw.giga-byte.com/Motherboard/Support/BIOS/BIOS_GA-8S661FXM-F.htm

Compiled by Steve Smith

YOU GET PAID FOR THAT?

Paradise needs programmers. Konami Digital Entertainment, makers of the Castlevania series, Metal Gear Solid, and Dance Dance Revolution, needs programmers to tough it out in its Honolulu studio. If you have a C/C++ background and experience with BREW or Java, you might be qualified to write mobile phone games for the company. Another job listing with the same company is looking for people versed in 3D programming to work on game engines for PCs and consoles. The listing states that knowledge of Japanese is a plus, and you'll likely be more willing to learn the language once you hear about Honolulu's climate. According to *USA Today*'s weather almanac, the average high on these Pacific islands is in the 80s every single month of the year, with average lows almost always in the 70s. So you'll get to program games and ditch your overcoat for good. ▲

RAW DATA

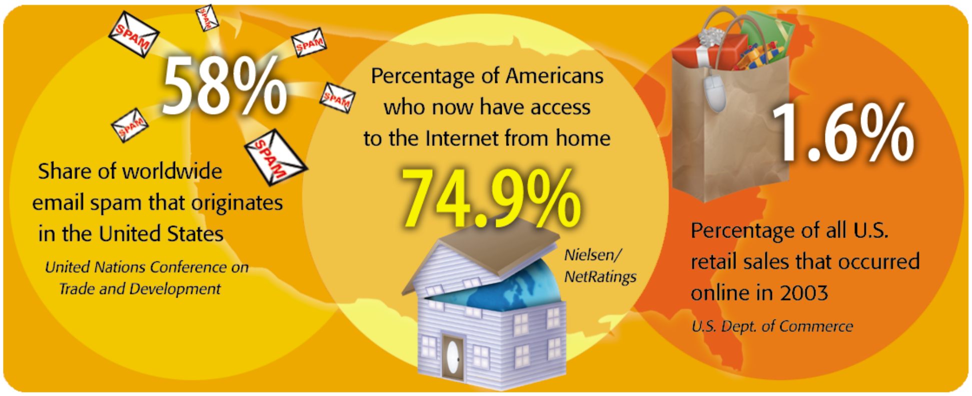

58%

Share of worldwide email spam that originates in the United States

United Nations Conference on Trade and Development

Percentage of Americans who now have access to the Internet from home

74.9%

Nielsen/ NetRatings

1.6%

Percentage of all U.S. retail sales that occurred online in 2003

U.S. Dept. of Commerce

Is File-Swapping Piracy Or Promotion?

Take that, RIAA! In the most methodologically rigorous tracking study to date of actual online file-swapping patterns and sales of specific music CDs, two university researchers say that, at best, P2P networks account for a very small fraction of the music industry's sharp decline. By tracking and comparing heavily downloaded songs with their retail sales, researchers at Harvard Business School and University of North Carolina found that in the case of CDs that sold over 600,000 units, P2P activity actually may help retail performance by selling an additional one CD for every 150 downloads of one of its songs. Professors Felix Oberholzer-Gee (Harvard) and Koleman Strumpf (UNC) concluded that in the worst case scenario, file swapping accounted for lost sales of only 2 million CDs in 2002, while as a whole, the industry sold 139 million fewer CDs between 2000 and 2002.

Half A Billion Phones & Counting

With cell phone sales passing the half billion-unit mark in 2003, mobile phones are the platform that just won't quit. The main beneficiary continues to be Finnish manufacturer Nokia, with more than a third of all worldwide handset sales, according to Gartner Dataquest.

Worldwide Mobile Handset Unit Sales In 2002 & 2003 (in thousands)

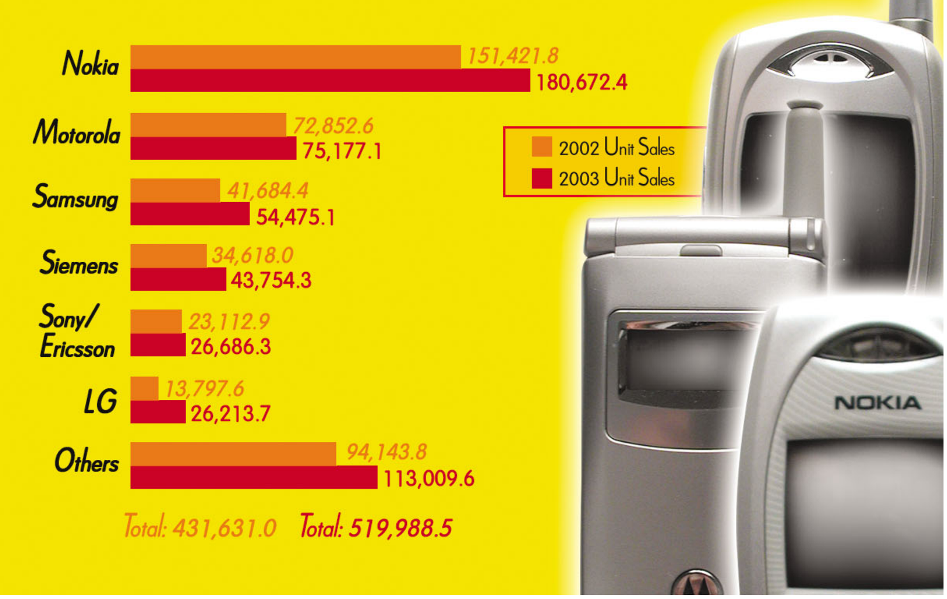

Source: Gartner Dataquest (March 2004)

Give *your* car the luxury of satellite navigation.

© 2003 Garmin Ltd. or its subsidiaries

The StreetPilot® III from Garmin has features found in expensive in-dash navigation systems, plus one very important quality the others don't have. StreetPilot III is portable, so you can have the best in GPS navigation in any car you drive. Look up addresses, attractions and other services, and StreetPilot III will create a route and provide turn-by-turn directions with voice prompts to get you there.

StreetPilot III and GPS V from Garmin. Navigation for every car.

GARMIN. www.garmin.com NASDAQ GRMN

Available at: BEST BUY COMPUSA

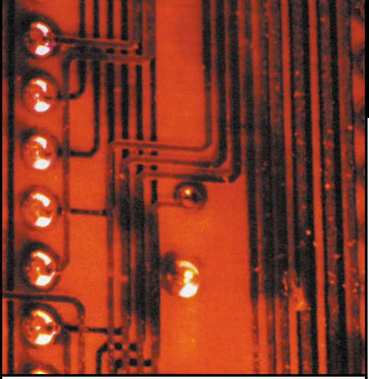

The Saint

by Alex St. John

Spyware 2.0

Alex St. John was one of the founding creators of Microsoft's DirectX technology. He is the subject of the book "Renegades Of The Empire" about the creation of DirectX and Chromeffects, an early effort by Microsoft to create a multimedia browser. Today Alex is President and CEO of WildTangent Inc., a technology company devoted to delivering CD-ROM quality entertainment content over the Web.

There's never a quiet day when you're running a software company. After months of dealing with the technical headaches generated for my company by Microsoft and Sun getting into a billion dollar tiff over Java IP, we're finally able to get back to the business of selling video games online. At least that's what I thought until one of the Internet's leading "removeware" products decided that my WildTangent Web Driver product was "spyware." My support manager informs me that support calls are rolling in as consumers naively accept the recommendations of their removeware product and delete the Web driver, thereby breaking all of the games they've been playing that depend on it. The standard Windows Uninstaller for our product works fine and doesn't leave a mess, but once a removeware vendor lists your products, consumers just assume your uninstaller can't be trusted.

> None of the removeware vendors appear to make any attempt to contact the companies whose software they are breaking . . .

It turns out that this particular removeware product is just generically deleting all files on the users hard drive under the folder name "WildTangent," then purging our registry entries, leaving nothing but broken games behind it. We're faced with the distinct possibility that some significant subset of the 78 million Web Driver gamers online are going to call us for customer support because they trusted this software and screwed up their computers with it.

Of course, I immediately set about getting in contact with this company to find out what's going on. It turns out that they are located in Europe and appear to be a company of about one guy. Attempts at phone and email contact are unsuccessful. "What exactly is it we're doing that makes them think we're spyware?" I ask, but the answer is troubling: "This guy says you are."

I call the CEOs of a couple other small online software publishers I know and ask if they are having this problem. The answer is a resounding "yes," all of them have a story about trying to reach one of these removeware vendors to try to get them to unlist their completely benign software products from automated deletion. Research reveals that they are mostly one- to four-person operations. Some are politically opposed to selling software for money and one says they make their software to drive "social change." They have no formal standards for differentiating legitimate software from malign software beyond sweeping generalizations that they apply selectively to some software products but not others. They also seem to rely on hysterical reports from consumers having troubles with their computers who ascribe all manner of nefarious behavior to various software products including our own. Many of the problems appear to be associated with users trying to delete programs manually on the assumption that the standard Windows uninstall process can't be trusted and then seeking help when they screw up their own computers.

None of the removeware vendors appears to make any attempt to contact the companies whose software they are breaking, inform them of the support requests they are getting from their users, or make any attempt to link to our support forums where some of the same issues are dealt with frequently and correctly. These companies presume all online distributed software guilty and malign.

To be perfectly fair, real spyware and adware has become a serious online plague, and these guys do a great job of battling it. They are also doing a great job of challenging legitimate software companies to look closely at their privacy and security practices to make sure that they are as clean and clear with consumers as they can possibly be. The problem is that in their zeal to stamp out evil software wherever it hides, they are also sewing consumer paranoia, damaging legitimate online software publishers, and empowering Microsoft to assume more control of the software that reaches your computer in the name of "security."

For my part, we're overhauling our newest generation of products to ensure that consumers have maximum understanding and control over everything they do. It's a lot of work and in the absence of any industry consensus on the subject we have to define our own standards for good privacy practice and hope the removeware guys agree with us. ■

Got a different POV? Email TheSaint@cpumag.com.

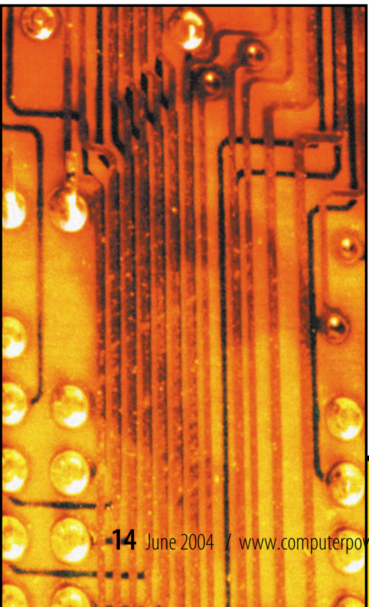

Dream Hardware

Couldn't we all use a little nonvolatility? If your dreams involve solid-state memory chips doing a little victory dance on a battlefield littered with smashed hard drives, get counseling. Or get a super-sized flash memory disk from M-Systems. Its space may not be infinite, but the Prelude MTS speakers are, sort of. And if your trust fund might as well be infinite, too, Mother says you should build the X-Wall. Er, right now.

by Marty Sems

M-Systems 90GB Fast Flash Disk 2.5-inch Ultra ATA

O, for the day when we can cast aside our space-heating hard drives in favor of purely solid-state storage. Good news, good news, Flash—that day is today. The capacity is now there, thanks to ultra-spacious nonvolatile flash drives such as M-Systems' (www.m-sys.com) 90GB example shown here. The only bubble in the bathtub is its $27,725 price tag. Of course, if you have massive space/speed needs and an expense account to match, consider the 2.5TB RAM disk that Dynamic Solutions International (www.dynamicsolutions.com) just sold the guv'ment. This OEM for Texas Memory Systems says that the RAM disk was actually a cluster of 40 RamSan-320 modules, each with 64GB of capacity. Retail would have been $4.8 million, but DSI tells us that the shadowy agency got the taxpayers a discount. Black Mesa Project, anyone? "It's all in flash, Gordon."

Infinity Prelude MTS Speakers

Great audio doesn't have to cost you six figures. Really, just a paltry 10 grand will do. Open your ears to the sound of Infinity (www.infinitysystems.com) with the Prelude MTS loudspeakers, which are aging very well in their fourth year. Ceramic-coated aluminum drivers keep every note, from the top end to the bottom registers, punchy and free of unwanted resonances. The 850W, 12-inch subwoofer also benefits from Infinity's Room Adaptive Bass Optimization System. Infinity offers two towers and a subwoofer as a $10,000 set, but it says that surround-sound lovers can add a $1,200 center unit and other speakers for 5.1 or 7.1 scenarios. But what's mere money, when Infinity awaits?

9X Media X-Wall

Live, in concert, it's . . . you. Start your own Zoo TV tour with your very own X-Wall. Yeah, that's X-Wall, as in video screens that cover the entire wall—like your ego! Visual whizzes 9X Media (www.9xmedia.com) build multipanel displays in a dizzying blitz of configurations. At the top end is a $294,000 16- x 10-foot edition that incorporates 49 Wide LCD Extreme Series monitors, each 30 inches corner to corner. In a sense, these united screens are one, but they're not the same. Even better than this real expensive thing, 9X Media says that 64-screen custom versions are available, as are larger LCDs and plasmas. Remember, nothing says you're trying to throw your arms around the world like a fisheye close-up of your nose writ large. Hanging Trabants not included. Achtung!

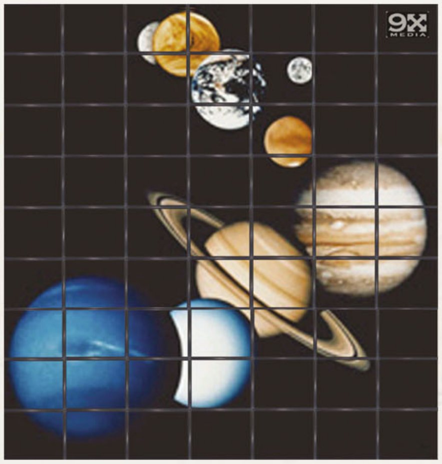

Extreme External Storage

Stockpiling From The Outside In

You don't just need ordinary external storage for your hordes of files. You need mountains of capacity you can access fast. Extreme storage can also mean sexy cosmetics and cool extra functionality. Lucky for us, our prayers for affordable, exceptional, mobile storage have been answered. This is a huge field, so we've selected some shining stars.

When testing external storage, your test platform doesn't count for much. Everything bottlenecks at the device's interface, whether it's FireWire or USB. We ran all products on a 3.4GHz Pentium 4 Northwood using an ASUS P4C800-E with 1GB of memory and a Radeon 9800 PRO. We used a 250GB Hitachi Deskstar to hold a 2GB folder of assorted files, wrote the folder to the target device, did a quick format of the Deskstar, read the folder from the external device, and copied it back to the Hitachi.

Hard Drives

Maxtor OneTouch Combo 250GB

We selected the 7,200rpm 250GB for its superior speed. Our benchmarks show there's often a disparity between synthetic and real-world tests, and in Sandra and HD Tach, the OneTouch looks laggardly. In PCMark and 2GB transfer tests, it emerges as arguably the fastest drive here.

Maxtor was the first major drive vendor to implement a shortcut button on its external drives. The feature has held fast throughout the shift to a sleek, anodized, aluminum chassis and a long button now backlit with two blue LEDs that flash to indicate drive activity. The package includes a rubber stand for vertical operation, but you can also stack it. This drive features FireWire and USB 2.0 ports with included cables for each type.

Setup is a snap. The drive automatically installs under Windows XP, but a wizard helps you customize the included Dantz Retrospect Backup Express software to run when you press the OneTouch button. Maxtor now lets you use the button to execute any app on your system.

The OneTouch Combo stays cool, runs extremely quietly, and even features a spin-down feature after one hour of inactivity to prolong its life span. In a long line of great external drives from Maxtor, this is the greatest yet.

SimpleTech 250GB SimpleDrive USB 2.0

It's as if SimpleTech's managers looked at the OneTouch drive and said, "OK, how can we make this simpler?" The 7,200rpm SimpleDrive keeps the aluminum chassis but omits the backup button. SimpleTech includes its own StorageSync app but seems to recognize that the idea of backup buttons is a bit silly.

More disturbing is that SimpleTech doesn't provide a means to help stack its external drives. Although it may sound trifling, we were a bit annoyed that the slip-on feet for vertical mounting left black marks on the chassis when removed.

OneTouch Combo 250GB
$350
Maxtor
www.maxtor.com
CPU CPU CPU CPU CP

This is a USB 2.0-only model, and it's performance is excellent thanks to its PATA Hitachi Deskstar drive, which helped the unit sweep all three of our synthetic apps for USB. Windows recognizes the drive automatically, and it ran flawlessly. An included Kensington lock slot to thwart theft is also noteworthy, although it doesn't entirely overcome that this drive is a bit overpriced.

250GB SimpleDrive USB 2.0
$365
SimpleTech
www.simpletech.com
CPU CPU CPU CP

Western Digital 250GB Media Center

WD's flagship 7,200rpm combo drive is most notable for its 8-in-1 flash card reader. WD places the traditionally rear-mounted power rocker in a front-mounted power button, with two backup buttons above it. The Automatic Backup button initiates the bundled Dantz Retrospect Express Wizard to help set up automatic backups. The Backup On Demand button performs selective backups on files and folders you specify through a different wizard. Neither button is reconfigurable.

The drive has two FireWire ports, plus two USB "B" ports, basically making it a mini-hub—probably its smartest feature.

CPU Ranking: 0 = Absolutely Worthless CPU CPU C 2.5 = Absolutely Average CPU CPU CPU CPU CPU 5 = Absolutely Perfect

This was the only drive we tested that required an outside driver for My Computer to recognize it. The drive can run over either interface, but the USB ports and flash reader require a USB cable. If you plug in both cables, the drive runs across FireWire, leaving other functions to USB. Taken together, this might explain the lagging performance in our synthetic USB scores. PCMark's inability to see *any* hard drive with the USB interface connected remains unexplained. Across FireWire, performance was excellent. Despite the highest price here, we admire the drive's functionality. It's the best unit here for mobile use.

(Subscribers can read our review of the Seagate 160GB External Hard Drive at www.cpumag.com/cpujun04/external.)

Out Of The Ordinary

IOGEAR BOSS (200GB)

The BOSS (Broadband Office Storage Server) is a concept so obvious it seems incredible that a relatively small company like IOGEAR capitalized on it first. In a nutshell, the BOSS is a NAS box that also serves as a VPN-compatible four-port 10/100 switch, a stateful packet-inspection firewall, and an FTP server. It's a 7,200rpm 200GB drive (a 120GB version is $100 less) stuck in a SOHO switch with some beefed-up firmware.

250GB Media Center

$400

Western Digital

www.wdc.com

CPU CPU CPU CP

Under the hood, the BOSS does a little of everything but lacks depth in some areas. For instance, the firewall misses such common advanced features as MAC address filtering and rule creation. The basics are there, though, and should serve SOHO and consumer needs. Configuration through the browser interface is simple. (There's no CD software.)

The painfully slow numbers in our transfer tests highlight the BOSS' chief weakness: Fast Ethernet is a poor interface for moving large amounts of data on a regular basis. However, for near-line data archiving and streaming of multimedia files across a LAN, the BOSS is fine, if only because it avoids the system resource bottlenecks inherent in USB/FireWire drives. IOGEAR integrated many complementary functions in one box, but we'd like to see the features deepened and the performance raised.

BOSS (200GB)

$500

IOGEAR

www.iogear.com

CPU CPU CPU CP

Iomega REV 35GB/90GB USB 2.0 External Drive

Iomega is still pounding away at removable cartridge products. Its latest is the REV, which uses a 35GB native (90GB compressed) 2.5-inch hard drive platter encased in a tough but lightweight, 3-inch square cartridge. Despite our aversion to proprietary product platforms, we admire the REV, which also comes in an internal ATAPI model. Read/write performance is quite good for a mobile backup device, and at $59 per single-pack cartridge, the price per gig is reasonable.

REV 35GB/90GB USB 2.0 External Drive

$400

Iomega

www.iomega.com

CPU CPU CPU CPU

Iomega uses its own REV UDF file format, which is why the drive appears in My Computer as a removable device. (This is also why our synthetic benchmarks don't see the drive.) REV UDF doesn't support NTFS-like file encryption, but it does offer password-based access protection. Additionally, you can make REV disks bootable, and Iomega tosses in software for file backup and disaster recovery. The bundled Norton Ghost for REV is an excellent add-on because most external drives don't provide for the simpler and more comprehensive approach of restoring a drive image.

Iomega markets the REV as a tape-drive killer for businesses, but if you can see yourself outgrowing a couple of traditional external hard drives or find the cartridge's size more convenient than full-size drives, you might be able to look past the REV's steep buy-in price.

Jukeboxes

Archos Gmini 220

At 6 ounces and about the size of a floppy, the Gmini 220 sports ruggedized rubber edges and a hardier feel than the iPod. The unit's only major shortcoming is its low-contrast, grayscale Photo Viewer screen, which painfully reminds us of old PalmPilot displays.

After about five minutes, we felt fluent with the Gmini's hierarchy-based navigation and interface. The included Type I/II CompactFlash slot is a huge plus. The device can record voice files with the integrated mic at up to 112Kbps and other audio via the bundled SPDIF/analog pigtail at up to 192Kbps. Playback quality in MP3 and WMA is quite good once you ditch the included earbuds, and battery runtime hovers around eight hours.

CPU Ranking: 0 = Absolutely Worthless CPU CPU C 2.5 = Absolutely Average CPU CPU CPU CPU CPU 5 = Absolutely Perfect

Outside The Box: How Our External Units Tested

Testing this wide a range of external storage products isn't exactly a scientific process, but we did it just the same. The following gives you an idea of what each unit can do.

					2GB folder transfer		Sandra 2004	HD Tach 2.7 (MBps)		PCMark04
Model	**Capacity**	**Rotation Speed**	**Buffer Size**	**Interface**	**Write**	**Read**	**File System Drive Index (KBps)**	**Avg. Write**	**Avg. Read**	**HD Test Suite**
Archos Gmini 220	20GB	4,200	2MB	N/A						
				USB 2.0	5:13	2:56	12049	N/A	16.4	failed
Creative Labs Nomad Jukebox Zen Xtra	40GB	4,200	16MB	N/A						
				USB 2.0	34:26	8:05	N/A	N/A	N/A	N/A
IOGEAR BOSS	200GB	7,200	8MB	RJ45	27:09	27:46	N/A	N/A	N/A	N/A
Iomega REV	35GB	4,200	2MB	N/A						
				USB 2.0	2:51	4:26	N/A	N/A	N/A	N/A
Maxtor OneTouch Combo	250GB	7,200	8MB	FireWire	1:54	2:03	17877	13.6	19.9	3513
				USB 2.0	2:15	2:35	21682	17.5	23.6	3680
Seagate External Hard Drive	160GB	7,200	8MB	FireWire	2:06	2:10	25575	17.8	23.8	3748
				USB 2.0	2:33	3:21	20178	17.2	18.8	3470
SimpleTech SimpleDrive USB 2.0	250GB	7,200	8MB	N/A						
				USB 2.0	2:17	2:33	22024	17.3	25.5	4149
Western Digital Media Center	250GB	7,200	8MB	FireWire	2:20	2:29	23110	19.5	34.4	3944
				USB 2.0	2:15	2:28	16825	14.4	17.1	Didn't detect

Kudos to Archos for having its jukebox appear in Windows as a hard drive rather than a proprietary removable device. Considering that the Gmini uses a 1.8-inch Hitachi Travelstar, our 2GB transfer tests were incredibly fast, and read speeds stayed within hailing distance of desktop performance. The Gmini lacks the iPod's games and PIM features, but we never use that stuff anyway. For audio apps and data storage, the Gmini 220 rocks.

Creative Nomad Jukebox Zen Xtra (40GB)

First, the bad news: After extensive technical support calls with Creative, it seems our Zen Xtra review unit had a slightly defective USB 2.0 interface. Our 2GB write test took more than a half hour, but Creative says that its in-house units can do this in about seven minutes. A replacement currently en route should help confirm this.

Gmini 220
$299
Archos
www.archos.com

Now, all the good news: With a sterling 98dB SNA, all the EAX effects you could ask for, and plenty of volume overhead, this unit sounds fantastic. In fact, it's the best-sounding player we've heard. The blue-backlit display shows more info than preceding Zen models, and navigation, while not on par with Apple, is pretty intuitive, especially given its excellent jog dial design. The same file system that kept us from running synthetic benchmarks lets the Xtra do searches from your palm, not just the desktop. Those with thousands of music files will love this.

While larger than it should be, you can't complain given the storage capacity for the price. For $300, Apple delivers 15GB with no data storage capability. (Another $100 gets you the 60GB Xtra.) The Zen Xtra also lets you replace its battery. Creative's MediaSource software is much improved, and the new Explorer-based Nomad Manager is great for dragging and dropping.

The Zen Xtra needs to be smaller, it's missing an FM radio option, and an SD slot would help with on-the-road transfers. The unit isn't enough to make us ditch our iPods, but if we were starting over, the Xtra probably packs the best bang for the buck of any jukebox available.

Last Words

Years ago, external storage was essentially for businesses. Now, options are hitting affordable prices. With software adding gigabytes *per day* to our file collections, we need cheap, fast, mobile storage. We've only looked at the high-capacity products based on hard drives. For moderate capacity storage in a much smaller design, subscribers can read our online reviews of flash-based products, plus our "Sound Of Storage" and "Tweak Guide" sidebars at www.cpumag.com/cpujun04/external. ▲

by William Van Winkle

Jukebox Zen Xtra (40GB)
$300
Creative Nomad
www.nomadworld.com

CPU Ranking: 0 = Absolutely Worthless 2.5 = Absolutely Average 5 = Absolutely Perfect

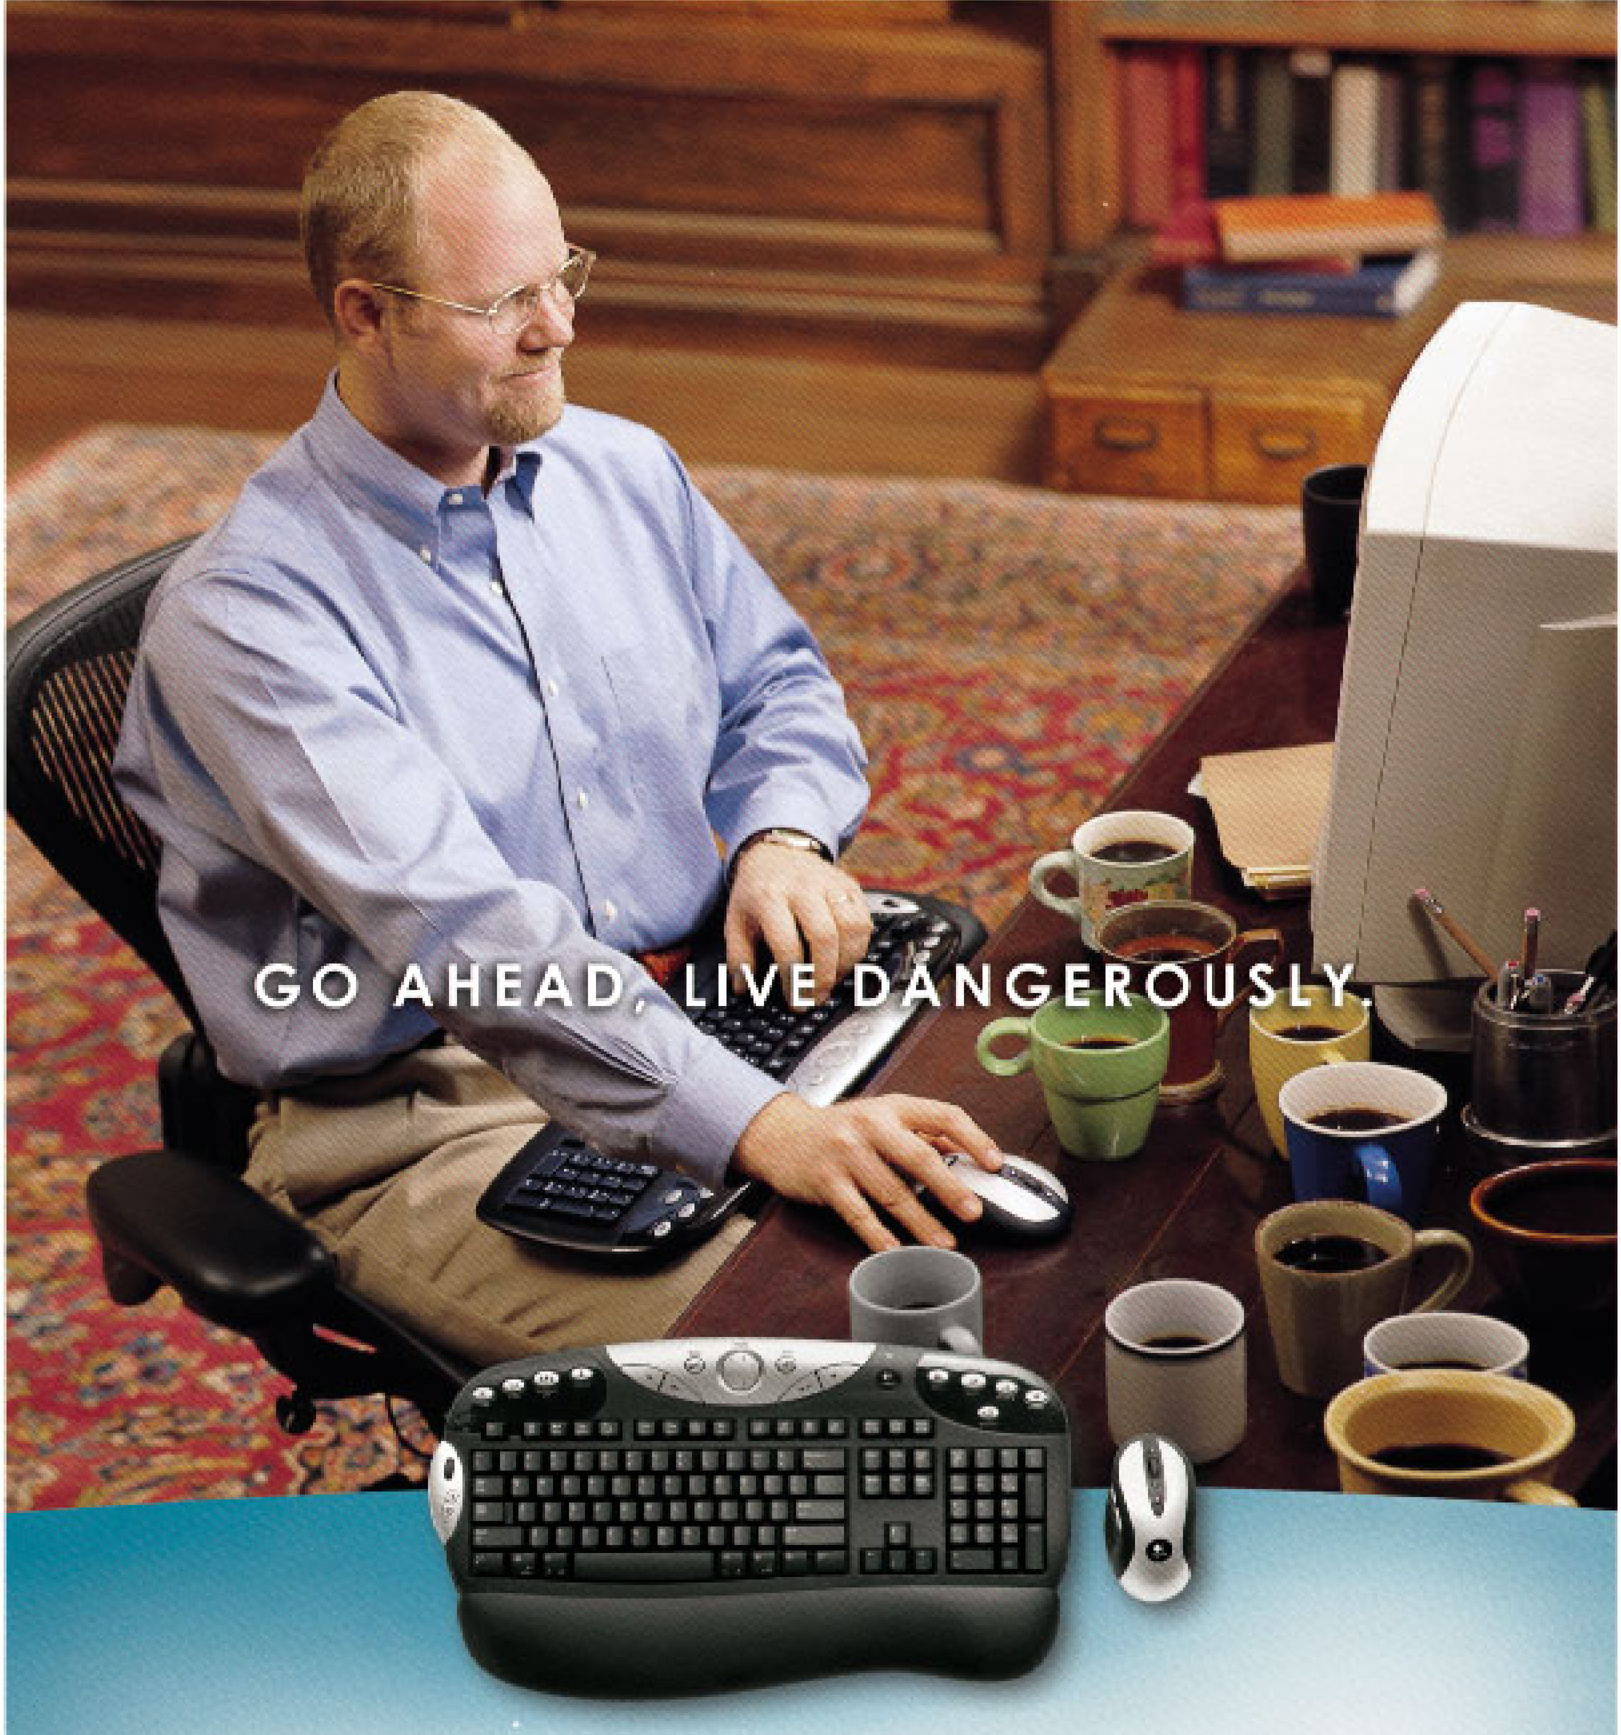

Logitech® Cordless MX Duo™. Announcing the end of tangled mice, stuck cursors and spilled coffee. With no cords to tie you up, slow you down, or hold you back, this sleek, one-touch keyboard and advanced optical mouse give you the control and flexibility to work your way – whatever the obstacles. What will you do with all that freedom?

www.logitech.com

Logitech

©2003 Logitech

Radeon X800 XT (R420)
ATI
Price To Be Announced
(905) 882-2600
www.ati.com

CPU **First Look**
PREVIEW

ATI Radeon X800 XT (R420)

With the 6800 Ultra, NVIDIA has really catapulted itself back up to the top and over ATI's Radeon 9800 XT by quite some margin. The pressure is now on ATI to come back, and based upon a recent trip to the company's HQ in Toronto, we can reveal that the follow-up product means business. Boards with the new R420 chip, based on TSMC's low-K 0.13-micron process, are only weeks away from shipping, according to ATI. The chip may not have NV40's 222 million trannies, but it's still got plenty at 160 million. Considering that each one has three nodes and a gate, that's over 500 million individual contacts. Cock up just one, and it's bon voyage into the rubbish.

The higher-end R420 part, the X800 XT, will pump along with 16 pixel pipes and six vertex shader units and is technologically well-matched to NVIDIA's 6800 Ultra. Some of the R420's fancy capabilities include being able to perform 80 concurrent shader operations per clock and 192 gigaFLOPS (a billion floating-point operations per second), all with a massive 8.8 gigapixel-per-second fill rate, which is becoming ever less important as we move to normal maps and such. When compared to NV40, somewhat glaring differences are the lack of support for FP32 and Shader model 3.0. However, ATI already supports some Shader model 3.0 features and with SmartShader HD, instruction capabilities are bumped up to 1,536 instructions (512 each of vector, scalar, and texture operations) per pass from 160, thanks to a new and improved shader compiler. Bear in mind that even a next-generation-looking game such as Half-Life 2 only has 30 shader instructions per pass.

The pixel shader enhancements include a more efficient F-buffer and multipass capabilities, increased temporal registers up from 12 to 32, and the addition of a facing register. Paralellism has always been important to ATI, and now with twice the vertex shader pipes of R3xx, the enhanced vertex processing capabilities (now with branching and looping) are beyond 750 million vertices per second. Start counting, or better still, start wishing for more detailed 3D models and scene animating effects such as soft shadows and blend shapes. And those six vertex shader units can perform up to 12 vertex shader operations per clock cycle.

Performance is one thing, but upping the ante with image quality is another. X800 features 3Dc (3D compression technology), a 4:1 compression for normal maps, which gives better storage capacity during execution of those shader instructions. Displacement effects created using standard resolution normal maps look great, but the Serious Sam 2 demo we saw had higher resolution 3Dc-enabled maps that provided more detail. Antialiasing is expected, and we're now at a level where 4X FSAA done with multisampling isn't all that expensive on frame rate, but going to a higher sample rate still is. With temporal antialiasing and centroid AA, the X800 XT has an alternative method. The multisampling techniques of R300 are carried over with gamma correct edge blending.

The X800 XT will use GDDR3 memory on a quad-channel interleaved GDDR3 256-bit memory interface. Currently, memory and core are aggressively clocked at or in excess of 500MHz (effectively 1GHz for memory), but final speeds are not yet decided. Unlike NVIDIA's top-end part, the ATI board boasts a single slot solution and actually uses less power than its predecessor, which is no mean feat. So no silly PSUs will be required, and it'll keep the SFF crowd happy. Even better news could be that the other R420-based board, called X800 Pro, will have a 475MHz core and 900MHz memory with 12 pixel pipelines and will be considerably cheaper.

The benchmarks ATI showed us had the X800 XT pretty much in a tie for the NV40 Ultra. Shipping dates for the X800 Pro and X800 XT are May 4 and May 21, respectively, which is before NVIDIA's next generation is due to ship. The battle lines are drawn. We'll finish off this piece by adding a disclaimer to remind you that everything we have seen was based on demonstrations run at ATI's HQ and not on our own test platforms (yet). But it's exciting times right now in the 3D graphics market with neither of the two big companies attempting to bring a knife to a gunfight. ▲

by Alex "Sharky" Ross

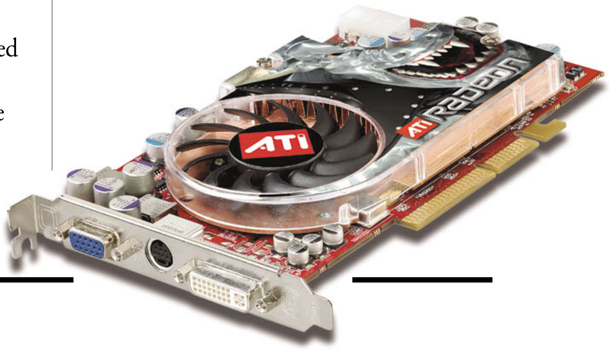

AMD Athlon 64 FX-53

Owning the crown at the super-high-end-CPU price point doesn't mean much financially because so few of those processors actually get sold. It's much more of a "bragging rights" issue. But if you believe in trickle-down economics via branding, being king of the hill indirectly helps sell the more affordable midrange processors. Intel and AMD have previously gone toe to toe with their Pentium 4 Extreme Edition 3.2GHz and Athlon 64 FX-51, respectively. Just recently, Intel released a follow-up P4EE 3.4GHz, so it was only natural for AMD to follow suit and crank the gigahertz knob a little to come out with the FX-53 processor.

Architecturally, the FX-53 is exactly the same as the FX-51, being based on the Opteron and packed with the usual goodies, including an integrated 128-bit memory controller, single HyperTransport Link, 128KB of L1 cache, and a larger (1MB) of L2 cache. The only difference is that the FX-53 runs at 2.4GHz (only a 9.5% boost over the FX-51). The process remains unchanged and is still based upon 0.13 microns, with a 193mm2 die size and 106 million transistors.

We pitted the top dog Intel Pentium 4 3.4GHz and EE 3.4GHz processors against AMD's finest. As you can see from the benchmarks, it's a close call, but the FX-53 seems to win the day where it counts. Under BAPCo's new SYSmark 2004 benchmark suite (a vast improvement over the previous version), the pendulum did swing in favor of Intel's Pentium 4s, which did not correlate to the rest of the benchmarks run. The same was true of Futuremark's PCMark 2004 CPU Tests, where again the Athlon 64 seems to strangely under-perform. However, the FX-53 met most game benchmarks we threw at it with a win. The Pentium 4EE 3.4GHz came close on some occasions, and although extremely fast, no longer holds the title for the fastest gaming CPU. Only with Quake III engine-based games does the Pentium 4 sneak ahead, but that engine is getting long in the tooth, so, unless you're still playing Jedi Knight II or Call of Duty, you'll most likely not care.

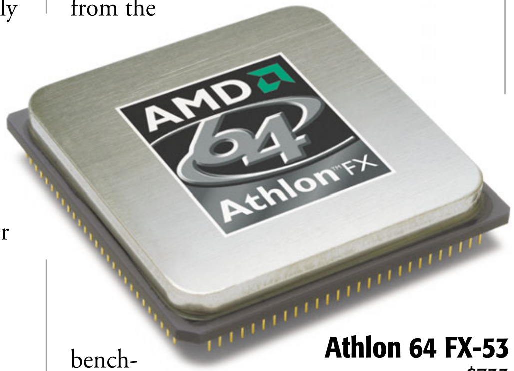

Athlon 64 FX-53
$733
AMD
(408) 749-4260
www.amd.com
CPU CPU CPU

So, while the FX-53's performance is, in most cases, second to none, it becomes difficult to recommend to anyone but the most hardcore gamer, given its excruciatingly expensive $733 price point (even though it is cheaper than the P4EE 3.4GHz). It's worth mentioning that when you consider the socket 940 and its necessity for buffered DIMMs as a platform, it's a dying breed. AMD's road maps clearly show the way ahead on the desktop being the Socket 939 and support for unbuffered memory (due out shortly), so the FX-53 isn't on most enthusiasts' lists.

Even the most ardent gamer will likely shy away from spending so much extra money for such little gain from either the P4EEs or FX processors. Much better value can be had with the ordinary Athlon 64, especially the 3400+, which is currently the best CPU buy in town. If you own a standard Athlon 64, pat yourself on the back and spend the money later on this year on some PCI Express and 3D graphics goodies instead; you'll thank us later.

Having said that, the FX-53's not quite done yet. Sure we got some fantastic scores using the older nForce3 150 Pro chipset, but there could be a little more juice left in the tank when using NVIDIA's nForce3 250GB platform. There are already rumors of an FX-55 before the end of the year, but we doubt that will sport the same pin count. ▲

by Alex "Sharky" Ross

Processor Benchmarks

	Athlon 64 FX-53	Athlon 64 FX-51	P4EE 3.4GHz	P4EE 3.2GHz
3DMark03 CPU Benchmark	793	748	786	735
Halo	59	57	58	57
Unreal Tournament 2003 Botmatch	107	102	92	88
Serious Sam: TSE	153	147	134	125
AquaMark 3	47	45	48	46
PCMark04				
CPU	4263	4231	5278	4851
Memory	5350	5364	5421	4483
SYSmark 2004 Overall	198	181	202	192
SYSmark 2004 Internet Content Creation	219	202	224	212
SYSmark 2004 3D Content Creation	211	194	214	204
SYSmark 2004 2D Content Creation	265	246	278	262
SYSmark 2004 Web Publication	187	174	194	182
SYSmark 2004 Office Productivity	179	161	183	175
SYSmark 2004 Communications	145	138	143	140
SYSmark Document Creation	207	185	215	205
SYSmark Data Analysis	193	165	197	192

CPU Ranking: 0 = Absolutely Worthless 2.5 = Absolutely Average 5 = Absolutely Perfect

Intel Pentium 4E 3.4GHz

Pentium 4E 3.4GHz
$417
Intel
(800) 538-3373
www.intel.com

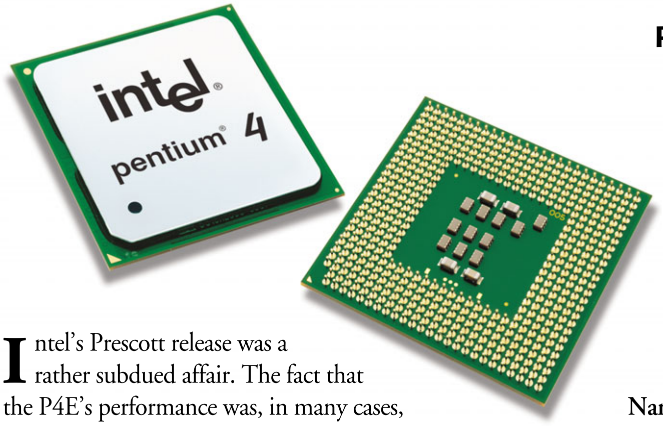

Intel's Prescott release was a rather subdued affair. The fact that the P4E's performance was, in many cases, lower than that of an equivalently clocked Northwood-based P4 didn't help set the world on fire. In fact, P4s based on Prescott cores haven't been flying off the shelves for an additional reason: the limited supply.

Intel's 90nm process is a new one and perhaps will be ramped up and/or better received in a socket 775 (Grantsdale) pairing or maybe later on once the rumored 64-bit support is added. So whether the launch of this faster clocked Prescott-based Pentium 4 3.4GHz will go the same way as the prior 3.2GHz and 3GHz processors remains to be seen.

Naming know-how. Incidentally, this now means that there are three Pentium 4s clocked at 3.4GHz: the Gallatin core-based Extreme Edition, the Prescott-based "E" designation, and, of course, the plain old Northwood version. Confused yet? It gets better. Intel is now set to take a page out of AMD's book and go with a nomenclature that should help good ol' Uncle Albert get his processor decisions right, even if he has been buying a lot of AMD CPUs over the past couple of years.

So it'll be out with the pure MHz naming convention we've become accustomed to and in with the model rating system. This will make it easier for Intel to market different processors based upon their features, such as Hyper-Threading, FSB, cache, and clock speed. This way, in the future, when there are cases of the same processors running at the same frequency, the 3.4GHz scenario can be avoided. The higher the number (based upon 5xx for the Pentium 4), the better the processor and more fully featured it will be. Until that time, an "E" denotes Prescott CPUs with an 800MHz FSB and 1MB of L2 cache.

Performance. The P4E 3.4GHz is a power-hungry sucker, and Intel estimates its key power load indicator, Thermal Design Power, at 103W. This is considerably higher than the Northwood's 89W at 3.4GHz. For now, the Prescott to have will be one based upon the D0 stepping, such as the new 3.4GHz, which Intel says will be able to achieve clock speeds up to the 4GHz mark.

Those of you with Northwood P4s must be scratching your heads again wondering why your older processor outperforms the current one at the same clock speed (except in certain situations). Despite Prescott's larger cache size, SSE3 instruction set, more advanced data prefetching, and deeper pipeline, the Northwood P4 running on today's unoptimized software still prevails. Exceptions include using Pinnacle Studio for Video Editing, in which we noticed considerable performance gains when using a Prescott-based P4 over the Northwood.

For the rest of you, if it's a Prescott you're after, now is not the time. It's going to take a new platform with PCI Express and DDR2 memory coupled with a hike in clock speed before Prescott is to become much more attractive to enthusiasts. The relatively aggressive pricing ($417) sees Intel target the Athlon 64 instead of the FX-53. Even so, the P4 E 3.4GHz is no match for the Athlon 64 3400+ in terms of performance, especially in gaming. Gamers who don't want an Athlon 64 would still be better with the Northwood P4 over a Prescott for now. Only if your usage model heavily relies upon apps such as video encoding or speech recognition does the Prescott make sense over the Northwood. ▲

by Alex "Sharky" Ross

Processor Benchmarks	P4E 3.4GHz	P4 3.4GHz	P4E 3.2GHz	Athlon 64 3400+
3DMark03 CPU Benchmark	708	707	675	712
Halo	52	53	50	55
Unreal Tournament 2003 Botmatch	81	83	75	99
Serious Sam: TSE	134	137	130	130
AquaMark 3	43	44	43	44
PCMark04				
CPU	5203	5231	4987	4312
Memory	5188	4789	5036	4003
SYSmark 2004 Overall	**194**	**191**	**190**	**182**
SYSmark 2004 Internet Content Creation	219	215	212	201
SYSmark 2004 3D Content Creation	212	208	202	294
SYSmark 2004 2D Content Creation	267	270	258	247
SYSmark 2004 Web Publication	194	187	175	171
SYSmark 2004 Office Productivity	178	171	172	165
SYSmark 2004 Communications	140	138	138	137
SYSmark Document Creation	200	196	188	187
SYSmark Data Analysis	202	197	195	170

CPU Ranking: 0 = Absolutely Worthless 2.5 = Absolutely Average 5 = Absolutely Perfect

Jetway MiniQ 860TWIN

Whether the usurper is your friend, significant other, child, or co-worker, it seems like there's always someone waiting to hop on your PC and steal your screen time. Finally, there's a workaround to this age-old problem for Windows users courtesy of Jetway: MagicTwin.

MagicTwin is a software bundle that essentially lets two users run two Desktops simultaneously off of one WinXP machine. The MagicTwin software hooks into a proprietary chip mounted on the motherboard. Jetway includes splitters for the PS/2 mouse and keyboard ports, and the user supplies a dual-head graphics card and additional sound card to complement the integrated motherboard audio. With this and the additional set of peripherals (mouse, keyboard, monitor, and speakers) installed on top of a configured Jetway MiniQ MagicTwin box, you essentially get two PCs for the price of a little more than one.

Considering this is a first-gen implementation for the mass market, the result is surprisingly good. A user on the left can watch a DVD while the user on the right plays Unreal Tournament. Both can be online concurrently (the IP address is shared) and both have access to the same applications, and it really does feel like you're running two separate systems . . . usually.

The MiniQ SFF

We loaded our 865G-based 860TWIN with 2 x 256MB of Corsair PC3200 TwinX, a Crucial Radeon 9800 Pro 256MB, a 250MB Maxtor DiamondMax 9 SATA drive, and a 16X DVD-ROM drive. We were a bit vexed to be forced into using a 3.0GHz P4 Northwood, as both our 3.4GHz Prescott and 3.2GHz Northwood chips would only run at 2.8GHz. Jetway says a forthcoming revision will fix this.

After playing with the unit for a while, we wondered if the speed limitation was enforced because of thermal issues. After 15 minutes of CPU burn-in under SiSoft SANDRA at 98% utilization, the MiniQ's processor temperature reported at 36 degrees Celsius—the same temperature it reported after sitting idle for an hour. There are no bottom air intakes, and the two side grilles are relatively small for a modern P4 design. Jetway uses a side-mounted fan on the CPU heatsink designed to pull air across the board and out into the 230W PSU's intake. However, this airflow generally misses the GPU on the left and the power circuitry on the right, neither of which have any nearby ventilation.

On the flip side, the MiniQ offers an IR remote control with integrated trackball and plenty of front and rear ports, including FireWire and optical SPDIF. The included data cables are sliced, save for the short SATA cable. Cable management through the chassis is decent, and the overall box appearance is attractive provided you don't get close enough to see how flimsy the bay and port doors are.

MiniQ 860TWIN
$299
Jetway
(510) 683-6588
www.jetwaycomputer.com

Benchmark Snapshot

3DMark 2003	5758
PCMark 2004	4659
SANDRA Memory test (average)	4597
AquaMark 3 (1,024, off/off)	45.22
AquaMark 3 (1,280, 4X/8X)	27.82
Comanche 4 (1,024, audio disabled)	57.54
Quake III Arena (640)	312.3
Quake III Arena (1,024)	304.7

Performance

Jetway recommends installing all apps and running any benchmarks in single-user mode with MagicTwin disabled. We did this and obtained a few test scores (with all devices, such as audio and LAN connectivity, enabled) just to confirm that the unit performs on par with other Springdale SFFs. Overall, the system performs solidly, and we had no unexpected benchmark glitches.

The thing to remember is that no matter how good the illusion is of having two systems, you're still stuck with the usual bottlenecks of a single CPU box. With Quake III running on one side, we ran through a series of applications on the other and found that hiccups only occurred during times of peak drive access, especially loading an application. Even DVD movie playback remained very solid until CPU utilization passed about 75%, whereupon dropped frames became increasingly common as CPU utilization increased. Again, this is roughly what you would expect in a single-user scenario. The only hitch we observed was when running a game on each side with CPU utilization pushed over 50%; one display would suddenly turn much darker than the other. Note, though, that Jetway expressly recommends running NVIDIA-based cards, which we didn't.

MagicTwin systems are best for people who could benefit from two side-by-side mainstream PCs and want to do it for a rock-bottom price. Given this, we can overlook the MiniQ's design shortcomings and recommend it as a very innovative product that won't violate your Windows EULA. (One license to one machine, right?) For those who want more flexibility, you can purchase bare MagicTwin-enabled boards from Jetway or its U.S. marketing partner, Mach Speed (www.machspeed.com). ▲

by William Van Winkle

CPU Ranking: 0 = Absolutely Worthless 2.5 = Absolutely Average 5 = Absolutely Perfect

MSI MEGA 180

MEGA 180
$299
MSI
(626) 913-0828
www.msicomputer.com

Probably too long overdue for its own good, MSI's MEGA 180 is still a worthy follow-up to last year's MEGA 651. The 180 fixes a few major shortcomings in its predecessor. For starters, it's much more attractive. The 180's all-black exterior is a much better choice than silver with quasi-wood tone orange, the new model's control buttons are more smartly arranged, and the revised red/green LCD is more useful and less annoying. The integrated 6-in-1 flash card reader and single external 5.25-inch bay remain unchanged, but side vents have been enlarged for better air flow, the remote control is better, and the instant-on FM radio and CD/MP3 player remain great.

That said, the 180 is still cramped. Sliding in the dual-fan, lateral blowing CPU heatsink is such a tight affair that the power wires on the back of the hard drive squeak as you push past them. One big oversight is that there aren't enough power leads to go around if you plug in an AGP card. The 250W Fortron Source PSU (nice choice) puts a SATA plug on one line, but there is no SATA support on the motherboard, and MSI provides no adapters or splitters. If you want to install an AGP card, the job is almost impossible unless you snip MSI's cable bundling. The documentation ignores AGP upgrading, but let's face it: You're not going to put an AGP card in a box like this. The nForce 2 (Crush18G) chipset only supports up to Athlon XP 3000+ chips and DDR333. If you want AGP graphics, you'll buy MSI's Springdale-based MEGA 865 box (which wasn't yet available as of this review). For the kinds of graphics needs you'd expect to have in a living room or bedroom setting, the integrated NVIDIA MX440 video should suffice.

That said, we are not sure the 180 is ready for prime time yet. Our first unit died on us just after driver installation. The second unit stayed up and running under Windows XP, but we gave up on AquaMark testing after 30 minutes. Comanche 4 stalled, and PCMark 2004 triggered a reboot. Looks like another "please refresh my BIOS" scenario.

We hope MSI fixes the bugs by the time you read this. Between the integrated Wi-Fi (via mini-PCI, just like a notebook), optical SPDIF in and out, four- and six-wire 1394, included TV tuner, dual-head output, and other worthy features, this is a great unit that just needs some improved stability. ▲

by William Van Winkle

Sony Cyber-shot DSC-F828

Cyber-shot DSC-F828
$1,200
Sony
(877) 865-7669
(408) 432-1600
www.sonystyle.com

Finally, 8-megapixel consumer-level digital cameras are flooding store shelves. The Sony Cyber-shot DSC-F828, the successor to the F717, was among the first. Sony gave the camera a complete makeover, fitting it with a more professional-looking black body and an improved positioning of controls.

The F828's list of features includes a Carl Zeiss f2.0 to f2.8 lens with a focal length of 7.1mm to 51mm (28mm to 200mm equivalent with a 35mm camera), 7X optical zoom, Hi-Speed USB 2.0 connectivity, PictBridge compatibility, full manual mode, and a NightFraming mode that lets you take pictures in even the darkest conditions.

The F828 supports Memory Stick, Memory Stick PRO, CF-I/II, and Microdrive cards. A backlit status LCD on top of the camera displays shooting modes and the InfoLithium battery's remaining life in both percentage and estimated number of minutes remaining.

Despite all of the impressive specs, the F828's image quality falls short of expectations. First, a quick-playback button would be better than twisting the mode dial to Play. Also, photos taken under fluorescent lighting on the Auto WB setting had a yellowish-pink tint. Chromatic aberrations and noise levels were the worst I've seen in a high-end digital camera. In Sony's defense, the noise likely arises from the 8MP CCD, which has to pack 3 million extra pixels on a sensor the same size as the F717's 5MP CCD. The smaller photosites likely struggle to capture the same amounts of light. Still, other 8MP cameras don't have such noise problems, suggesting that Sony may have rushed its 8MP technology to market too soon.

The F828 did many things well, such as producing outdoor photos with accurate, brilliant colors. Also, NightFraming shots turned out as well as those with the lights on. The F828 aspires to be a great camera, but Sony needs to address some of the F828's image-quality issues in the next model. ▲

by Kylee Dickey

CPU Ranking: 0 = Absolutely Worthless 2.5 = Absolutely Average 5 = Absolutely Perfect

Immaculate Wi-Fi Reception

The 802.11g USB 2.0 adapter with twist-and-shout flexibility!

IOGEAR's ultra fast Wireless-G 802.11g USB 2.0 adapter can help you connect to a wireless network with "higher power." It bends and twists to ensure optimal reception so you can work at up to 54Mbps at home or anywhere a hotspot is available.

Installation takes just moments. All IOGEAR Wireless-G products offer the latest security and are backward compatible with 802.11b.

So "retire the wire" without sacrificing Internet speeds - pick up IOGEAR's line of Wireless-G products today!

GWU513

Wireless-G to USB 2.0 Flex Adapter

- Uses USB port, not cardbus or PCI slot
- No configuration hassles
- Flexes and pivots for optimal reception

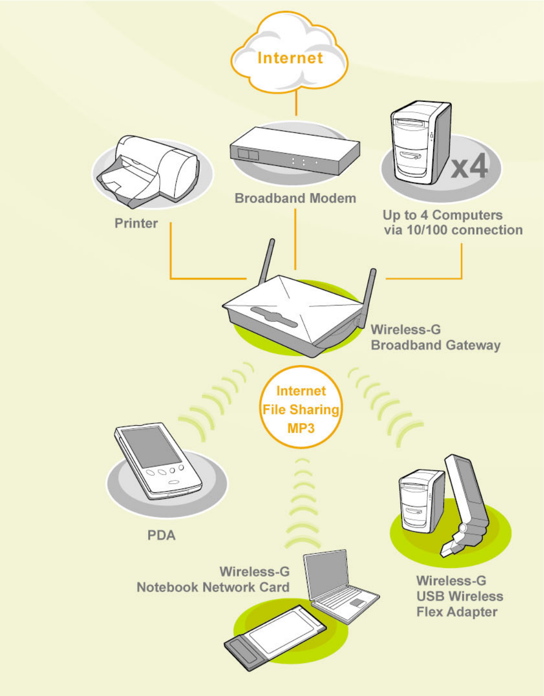
Internet
Printer
Broadband Modem
x4
Up to 4 Computers via 10/100 connection
Wireless-G Broadband Gateway
Internet File Sharing MP3
PDA
Wireless-G Notebook Network Card
Wireless-G USB Wireless Flex Adapter

GWA501

Wireless-G Broadband Gateway/Print Server

- Fast setup in just a few simple steps
- Includes everything for effortless installation
- Built-in firewall, supports WPA and WEP
- Built-in print server shares your printer with all connected computers

GWP511

Wireless-G Notebook Network Card

- Low-profile high-gain antenna achieves favorable sleek design
- Included Software Access Point allows you to turn your notebook into a wireless Access Point

Retire the Wire.

IOGEAR Inc.
23 Hubble Irvine, CA 92618 P: 949.453.8782 ext. 2011 http://smc.iogear.com

Tech Data

IOGEAR®
New Thinking, New Style

© 2004 IOGEAR, Inc. All rights reserved. All other trademarks are the property of their respective owners.

TDK Indi DVD 880N 8X/4X/12X DVD+RW 8X/2X/12X DVD-RW USB 2.0 FireWire 400

Indi DVD 880N 8X/4X/12X DVD+RW 8X/2X/12X DVD-RW USB 2.0 FireWire 400
$369.99
TDK
(800) 835-8273
www.tdk.com

TDK is still swinging in the DVD drive market, and it's going to blacken some eyes with 8X drives like the 880N. This external burner comes with both 400Mbps FireWire and High-Speed USB 2.0 connectors and cables. It supports Mt. Rainier and has a deep, dark secret I'll get to in a minute.

The 880N's curvy black shell looks fine until you mar it with fingerprints, which will stand out like bugs on a wedding cake. I recommend that you put it where you want it, give it a good wipedown with a cloth, and let it be. With its seriously nonskidding rubber feet, it will stay where you put it. TDK thoughtfully included right and left audio output RCA jacks on the back panel, along with a FireWire pass-through port for daisy chains.

Our test PC for this review had a 3.2GHz Intel P4, 1GB of Kingston PC3200 DDR, an ASUS P4P8X mainboard, and Windows XP Pro. The 880N gave us the fastest 4GB DVD+/-RW burn times I've yet seen, with 15:43 (minutes: seconds) for DVD-RW and 14:59 for DVD+RW. This leads me to that dark secret I mentioned: This Indi wrote a DVD-RW at 4X speed, even though the drive is only rated at 2X. Bonus!

The 880N's 12:45 DVD-R and 12:19 DVD+R burn times for a 4GB folder are faster than that of any drive I've reviewed, except for the Memorex True8X, Plextor PX-708A, and LG GSA-4082B. CD writes were just as speedy at 3:44 for a 700MB CD-R and 5:16 for a 639MB CD-RW. The drive's Nero DVD Speed results included competitive 201ms random and 319ms full-stroke access times, but the drive could only read DVDs at 4X across the board.

The mostly PC software bundle from Roxio also includes Toast Lite for the Mac, which is nice. However, the price isn't. The drive retails for $369.99, a jump up from the $329.99 price TDK predicted at CES. In light of $120 online prices for internal burners at similar speeds, $369.99 seems excessive for an external. Heck, even the new 12X ATAPI Plextor only costs $209. ▲

by Marty Sems

Epson Stylus Photo R200

Stylus Photo R200
$99
Epson
(800) 463-7766
(562) 981-3840
www.epson.com

With the R200, Epson introduced the most advanced, low-cost photo printer ever to hit store shelves. That's good news for cash-strapped digital camera enthusiasts, who may well find that the R200 more than satisfies their printing needs.

The R200 has extras you won't find on many printers in this price range, including a six-color print engine that uses individual ink cartridges—that's a huge step toward more affordable printing, as most so-called photo printers at the $100 price point use super-expensive multicolor cartridges. The R200 also has Print Image Matching technology and can print directly to the label side of printable CDs and DVDs.

At their fastest setting all Stylus Photo printers create spotted text, and the R200 was no exception. Default resolution text looked much clearer and darker, but in terms of quality, output was still not up to par with an average HP Deskjet. Printing speeds were on par with Deskjets, however, as the R200 produced times similar to HP's $130 5650.

The R200 creates glossies slowly compared to Canon's photo printers. I printed borderless 4- x 6-inch photos in about two minutes, and borderless letter-sized prints were done in five minutes.

Final prints looked nice, and the R200 produced image quality similar to more expensive Stylus Photo models, such as the R300. In some cases, though, I preferred both Canon and HP printouts to those from the R200, as this Epson's colors lacked the vividness and clarity I expected. Darker tones, such as black and purple, had excellent depth, but whites and yellows looked dingy in comparison to output from Canon's i-series photo printers.

In the end, that level of quality is fine for new digital printing enthusiasts. The R200 is an excellent way to get started on making nice prints at home for a minimal investment. ▲

by Nathan Chandler

CPU Ranking: 0 = Absolutely Worthless 2.5 = Absolutely Average 5 = Absolutely Perfect

Heatsink Hardball

Aerocool vs. Gigabyte vs. Zalman

In a matchup that looked to be a Goliath-vs.-David (and his brother) scenario, we pitted the titan of air cooling, Zalman, against relative heatsink upstarts Gigabyte and Aerocool. We tested on an ASUS P4C800-E board with a 3GHz Northwood Pentium 4 running a 98% burn-in load in SANDRA. Let's see how the Korean king of cooling fared against its foes.

Aerocool HT-101

The HT-101 drew two strikes from us almost immediately. First off, there is no adjustable fan control. The 80mm case fan that screws into either side of the heatsink spins at 2,500rpm, period. Second, installation is a far more time-intensive chore than with the two rival units. On our P4 reference platform, we had to uninstall the board's retention module and underside plate and then replace this with Aerocool's own plate and four rubber standoff tubes. All told, there are 10 screws in this unit. By contrast, Zalman uses two.

However, what the all-copper HT-101 lacks in finesse it makes up for in design. The heatsink is a 3-inch tall stack of 31 fins intersected by a tall U shape of three 6mm heatpipes. At default CPU speeds, the HT-101 ran the coolest of our three competitors. In fact, with a noise output of only 22dB, the HT-101 is arguably the most effective "silent" air cooler for standard run speeds. We also like the blue-tinged enclosure, blue LEDs in the fan, and UV sensitivity.

If you're an overclocker, though, this may not be the HSF for you. The variable speed fan is clearly an advantage under stressed situations.

Gigabyte PCU31-VH (3D Cooler-Ultra)

Gigabyte's heatsink is 3.6 inches tall and is just a pinch heavier than the Zalman. The unit is a stack of copper fins shot through with four copper heatpipes. A 2,000 to 4,500rpm fan sits in the stack's center, blowing out heat in every direction. In fact, we could feel the PCU31's breeze up to about 3 feet away, and when running at full speed (42dB), you can't help but think of using it as a hair drier. But at its lowest speed, the cooler is nearly silent.

Gigabyte tacks four, blue LEDs onto the cooler's top rim for eye candy. Installation is just as easy as with the Zalman, and Gigabyte supports the same three-processor platforms. Gigabyte also does the smart thing by making its fan controller-mountable in either a PCI backplane slot or a 3.5-inch drive bay, although it entails more cable clutter.

Zalman beats Gigabyte across the board except for tests in which we ran under a full CPU load with maximum fan speed. This should come as welcome news for OCers who care more about speed than noise.

Zalman CNPS7000A-Cu

The nearly all-copper 7000A-Cu is fashioned like a bowl of thin copper blades radiating from the center. Within the bowl rests a 92mm fan that whispers quietly at 25dB when cranked to its maximum of 2,400rpm. At the standard speed of 1,350rpm, the fan is almost silent at 20dB.

The 7000A-Cu comes with a dial-based fan speed controller that plugs into your mobo's CPU fan header and then connects at the back to the cooler's power connector. For bench testers and those who tend never to screw on their side panels, this is fine, but the lack of an external controller may bother some.

HT-101
$55
Aerocool
(510) 661-9948
www.aerocool.us

PCU31-VH (3D Cooler-Ultra)
$60
Gigabyte
(626) 854-9338
www.giga-byte.com

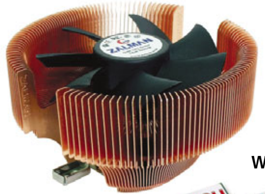

CNPS7000A-Cu
$50
Zalman
(714) 693-8880
www.zalman.co.kr/usa

The unit measures 4.3 inches in diameter, so some systems may have trouble accommodating this heatsink. Additionally, at 1.7 pounds, the 7000A-Cu is well out of weight spec for the Socket 478, 462, and 754 platforms it supports, so be careful moving a 7000A-Cu-equipped system.

Zalman's design is an excellent cooler considering its low noise. The 7000A-Cu realized a 2 to 4 degree Fahrenheit improvement over Intel's reference heatsink fan and topped out at 47 degrees in silent mode when overclocked under a full load. ▲

by William Van Winkle

Heatsink Head-To-Head Results

		Default				10% OC			
		Min RPM		Max RPM		Min RPM		Max RPM	
		Idle	Loaded	Idle	Loaded	Idle	Loaded	Idle	Loaded
Aerocool	HT-101	30/86	44/111	N/A	N/A	32/89	46/114	N/A	N/A
Gigabyte	PCU31-VH	33/91	47/116	31/87	40/104	34/93	49/120	32/89	41/105
Zalman	CPNS7000A-Cu	32/89	45/113	30/86	42/107	30/86	47/116	30/86	44/111

CPU Ranking: 0 = Absolutely Worthless 2.5 = Absolutely Average 5 = Absolutely Perfect

Anand's Corner

by Anand Lal Shimpi

GeForce 6—Finally

Anand Lal Shimpi has turned a fledgling personal page on GeoCities.com into one of the world's most visited and trusted PC hardware sites. Anand started his site in 1997 at just 14 years old and has since been featured in USA Today, *CBS' 48 Hours and* Fortune. *His site—www.anandtech.com—receives more than 55 million page views and is read by more than 2 million readers per month.*

I was talking to some game developers a while back about NVIDIA, and at the time, its flagship chip, known as NV35. I asked the developers what they thought about NV35 and they responded with a blunt: "We're surprised NVIDIA has managed to get away with what they have been for so long." Indeed NVIDIA did get away with a lot, although the NV3x series of chips were clearly inferior to ATI's R3xx family, NVIDIA honestly got off easy. It got a lot of flack from the community, especially when it came to that dirty word "optimization" and its drivers, but people still bought NV3x parts and NVIDIA was still alive and well by the end of the 12-month FX product cycle.

The company was a bit bruised, and any egos that were present had since been deflated, and it was almost fitting that NVIDIA should feel the pain it put ATI through so many product cycles in a row, where ATI's fastest chips just never seemed to be able to dethrone its competition. NVIDIA lost a lot of ground and credibility with the GeForce FX series of parts, and honestly, most NVIDIA owners in 2003 were still NVIDIA owners because they had yet to upgrade their GeForce4s—quite possibly the last good NVIDIA chip . . . until now.

> **. . . and it was almost fitting that NVIDIA should feel the pain it put ATI through so many product cycles in a row . . .**

Talk of NV40 (NVIDIA's next-generation architecture) started happening early, and it started off the same way as talk about NV38 did before it, and the same way as NV35 did before that: "We've learned from our mistakes," "shame on us," and so forth. The first time you hear it, you believe it, especially with a company with NVIDIA's once-strong reputation, but by the third time, you're unwilling to believe anything. Those three strikes put NVIDIA almost out of the credibility ballgame with NV30, NV35, and NV38, so I was dubious about NV40's "we've learned from our mistakes" message.

Eventually some of the architecture was revealed and it became clear that NVIDIA had fixed what it said it would. NV40, the first chip in the GeForce6 line (note the dropping of the FX from the name), has a number of improvements that improve performance a lot. NVIDIA has finally admitted that the NV38 design was really designed to render four pixels per clock, and in some cases capable of doing eight per clock. For a long time NVIDIA skirted this issue, but now that they are no longer bound by the NV3x architecture, it's easy to point out flaws (not so easy for the poor saps who bought NV3x cards). While NV3x could spit out between four to eight pixels per clock, NV40 is a pure 16 pixels per clock chip, and in some rare cases can render even more.

NVIDIA included more math units on the chip and removed a handful of limitations constraining sequences of instructions to attain maximum performance out of NV40. While the NV3x architecture relied very heavily on NVIDIA's run-time-compiler in the driver to run shaders at a reasonable speed, the NV40 architecture is inherently much more flexible.

NVIDIA has also added a dedicated video processor on the NV40 that is about as complex as the original GeForce. What can the dedicated video processor do? Finally offload CPU intensive video encoding from your CPU, onto a very fast and very expensive GPU. The result? Imagine real-time 1080i MPEG-4 or DivX encoding done on your computer, virtually regardless of your host CPU speed. The other benefit? NVIDIA says that the video processor will make it to all NV4x chips, so eventually there will be a $79 card with this functionality on it.

All this comes at a cost, the NV40 is bigger than Intel's largest desktop chip (at 220M transistors) and you can estimate an almost Itanium-like die size. It doesn't matter much to end users, but don't plan on overclocking NV40 much. Also, the company needs a smaller manufacturing process to bring NV40 to the mainstream.

Performance-wise, the first NV40-based card, the $500 GeForce 6800 Ultra, is fast. At more than 1,280 x 1,024, it's at least twice as fast as the Radeon 9800XT. ATI is about to respond with the Radeon X800, but ATI may break a sweat this time around. ■

Anand is listening at Anand@cpumag.com.

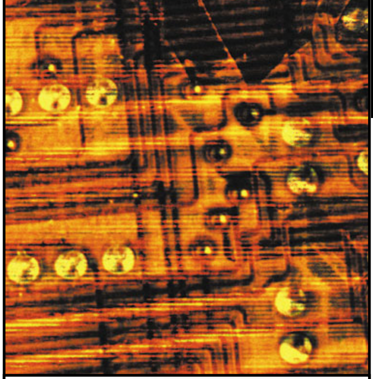

The Shark Tank

by Alex “Sharky” Ross

My Battery Life

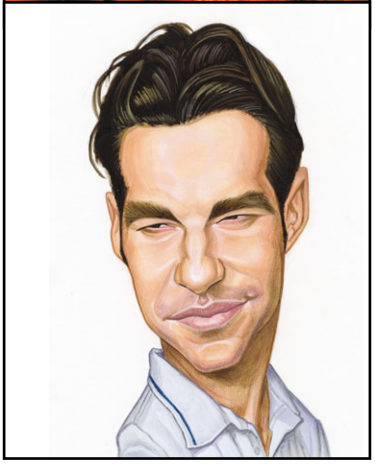

Disrupting Reuters’ newswire with a cheery Christmas greeting at age six, Alex “Sharky” Ross became an avid computer user/abuser, eventually founding popular hardware testing/review Web site SharkyExtreme.com. Exposing shoddy manufacturing practices and rubbish-spouting marketing weasels while championing innovative products, illuminating new technology, and pioneering real-world testing methods was just a front for playing with the best toys. The site acquired, he left in 2001. A London native and London School of Economics graduate, Alex currently swims in Silicon Valley.

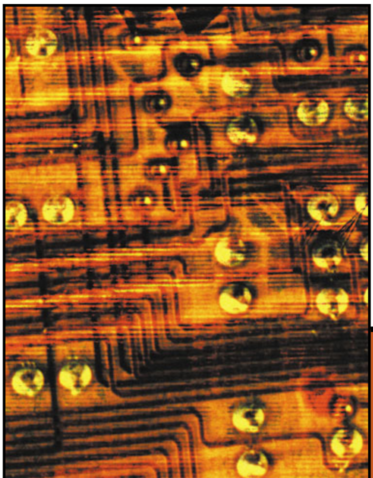

I’ve been using notebooks for years and years, back when they were heavier than four sacks of potatoes and about as zippy as a Citroen 2CV on 87 octane. Over the years things have improved, and despite many growing pains, I’ve stuck with them and learned to love and use them, despite being endowed with a world-class desktop PC at the same time.

Maybe it’s because I don’t actually want to be tied down to a desk when I work (except when I’m playing Far Cry, of course), or maybe I just keep thinking that mobile should indeed mean what it states? I’ve probably owned at least 10 notebooks (that I can remember), which is far less than my extensive desktop collection. Of course, that’s because notebooks are usually slow enough in terms of innovation and implementation of new technology that they rarely warrant any kind of upgrade (which also entails purchasing a whole new system, unlike on the desktop where parts can be interchanged—in my case on a daily basis). Yes, it’s a bloody expensive habit, too. That’s it. Some of the better changes have been larger screens, higher resolutions, cheaper prices, lighter weight, faster CPUs and hard drives, and better graphics, even to the point where I can have a decent game.

Using this notebook has made me feel mentally untethered.

But for the first time in a long time, I’m actually using a new notebook that *has* made a difference. Enter the Dell Inspiron 600M. It’s powerful (a 7018 3DMark score is nothing to sniff at), it’s light, and it has decent graphics and a lovely 14.1-inch SXGA+ screen with a native 1,400 x 1,050 resolution. The 600M is connected and truly mobile, sporting Intel’s built-in wireless 802.11b, but holy Batman, it’s been over seven hours since I’ve had the thing attached to a power supply plugged into the wall. I’ve gotten carried away for hours researching and writing my NV40 article, answering email, chatting online, and reading the latest football scores, all without noticing any of that dreaded heat dissipation on my lap. Finally, I’m mobile, and this time without that burning sensation!

I remember with some trepidation during an IDF last year that I was somewhat unconvinced by the Centrino platform. It didn’t seem to have the speeds of the Pentium 4, and I must have taken the improvements to battery life with a grain of salt and/or swept it under the carpet. Without getting into technicalities, from my own perspective, using this notebook has made me feel mentally untethered. Instead of gazing at the battery every so often, nervously typing away trying to get as much done before having to get to a socket for another charge, I am now able to work/use the notebook for virtually as long as I can bear to take it. It recharges when I do.

With a 1.6GHz Pentium M processor, 512MB of 266MHz DDR memory, a 40GB hard drive, and a 32MB ATI graphics chip, I’m in mobile happyland. The 48WHr (six-cell Li-Ion battery) additional modular battery provides dual battery support (especially worth the additional $89) and plugs into the DVD/CD-RW combo bay to bump battery life up from around five hours to almost eight hours. But wait, there’s more: For an extra $49, I can use its Bluetooth capabilities and communicate with my Sony Ericsson T616 mobile phone, when I figure out how to use it, that is. And for such a fully featured notebook, it only weighs in at 4.2 pounds, so you won’t break your back lugging the thing around. You can even dip into the 3-pound mark if you replace the media drive with a plastic dummy piece.

Having used the 600M notebook for a while now, there’s no chance that the manufacturer (Dell in this case) is ever going to get the thing back. Usually hardware sent in for review is sent back after the review process has run its course. But in this case, it’s taking a British Airways Business Class flight from SFO to London, where it will surely last almost the entire flight on a single charge. The only problem is, it’s going home with Sharky’s mom because she was the one who actually purchased it. Time for an upgrade, and time for me to start looking for a Centrino system of my own. . . . ■

Email some 1,400 x 1,050 images to sharky@cpumag.com.

by Kyle Bennett

Benchmarks Tell You What?

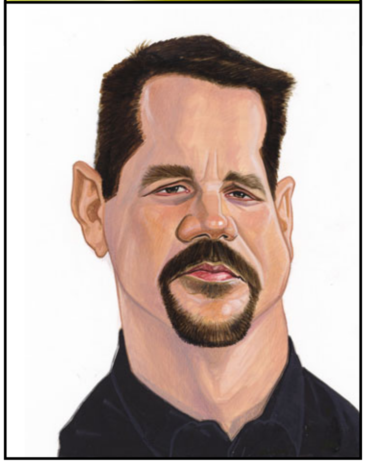

Kyle Bennett is editor-in-chief of HardOCP.com (hardocp.com), one of the largest and most outspoken PC-enthusiast sites on the Web. HardOCP.com is geared toward users with a passion for PCs and those who want to get cutting-edge performance from their systems. Beware, though, Kyle is known for his strong opinions and stating them in a no-nonsense manner while delivering some of the most in-depth reviews and PC hardware news on the 'Net.

As you read this, the video card wars have most likely been fully unleashed or are very close to it. NVIDIA's GeForce 6 series technology was introduced in April, and there's no doubt that it's truly a generational leap over the previous GeForce FX 5 series GPU. As of writing this (April 20), we have yet to get our hands on the new ATI Radeon X800 (page 20) series of VPU, but we have seen a great deal of information on it. The Radeon X800 family looks to be serious competition for the GeForce 6 family. One source at ATI even alluded to its 12-pipeline VPU (both ATI and NVIDIA's top-end parts have 16 pixel pipelines) being direct competition to NVIDIA's current 16-pipe flagship part. No matter what happens, we will soon all know the truth. Or will we?

There are probably hundreds of Web sites and print media outlets that have the facts, but is cold, hard benchmark data what the buyer needs? There have been more than a few instances in the last year that have truly made me question how we benchmark video cards. In fact, I have questioned the entire nature of video card benchmarking. I have lived and died by the benchmark for almost seven years now, and I came to question my own practices after I had done a video card review and was headed out to San Francisco to do a segment on the "Screen Savers" TV show, for which I was supposed to talk about the video card in question. I realized a couple of days before I left that I had never played a game using the video card, yet I had written the review. It was just like the cartoons with a light bulb pinging to life above my head. Here I am, supposedly some "expert" on this subject, and I had never done the one thing that millions of people buy the product for: to play 3D PC games. Combine this with the fact that pure cheating was being exposed in synthetic and canned benchmarks at about the same time, and I was left to do some thinking about what exactly HardOCP.com's goals are and what our duties are to our readers.

It may not be abundantly clear, but most of our readers are not hard-core technology enthusiasts by a long shot. They are people trying to make a real buying decision. If HardOCP.com is ignoring real-world gameplay, I think it can easily be argued that we are forsaking the exact people that we're trying to help. So, I pulled back and started trying to formulate exactly what we needed to do. Well, if people buy video cards to play games, then we should be showing exactly what real-world gameplay is all about. All of a sudden, another light bulb appeared in an otherwise dark room. HardOCP.com needs to stop using benchmarks. Benchmarks mean nothing to end users of the technology. Gameplay means everything. HardOCP needs to stop "reviewing hardware" and start "evaluating the gameplay" provided by the hardware.

HardOCP.com needs to stop using benchmarks. Benchmarks mean nothing to end users of the technology.

Thinking how to implement this soon became clear. Play games on the hardware and report back to our readers in an objective and subjective manner. We played tons of games, took the most video-intensive levels from those games, and used them for gameplay evaluation. We now use a graphing system that shows the frame rate every second through a handpicked section of the game in question. We also show these frame rates at what we think to be the best image quality supplied by the video card, in terms of resolution, antialiasing, and anisotropic filtering. We also describe our experience. We are using live, repeated run-throughs in games when we cannot use recorded demos. For games that allow demos, we don't release our demos publicly, in case interested video card companies might "optimize" to get better scores.

Please come read, and give us your feedback. I know one thing for sure: HardOCP.com is moving in the right direction by moving away from canned benchmarks and synthetics that don't in any way, shape, or form tell you, the person buying the card, what gameplay experience to expect. ■

Talk with Kyle at kyle@cpumag.com.

Lighten Up.

It wasn't so long ago that if you wanted to bring your data with you, your hardware had to come too. Well, that's all in the past now thanks to the new SanDisk Cruzer™ Mini. At under half an ounce, the Cruzer Mini goes anywhere, and with up to 256MB of flash memory, it takes a lot with it. Hi-speed USB 2.0 compliant, SanDisk's Cruzer Mini lets you transfer data faster than ever before. And at less than 8mm thick, all it needs is a single USB slot. Just drag, drop, and go. It's all about convenience – no extra hardware, no worries.

So whether you're installing utilities or updating print drivers for the umpteenth time, you can take your files with you wherever you go – and you won't even need your dolly.

For more information on SanDisk Cruzer Mini, visit www.sandisk.com/smart

HI-SPEED CERTIFIED USB
HI-SPEED USB 2.0

SanDisk®

SanDisk and the SanDisk logo are registered trademarks and Cruzer is a trademark of SanDisk Corporation. Other brands and products are trademarks of their respective holders. ©2004 SanDisk Corporation. All rights reserved.

PC Modder

Tips & Tutorials

Modding does the body good. A PC's body anyway, inside and out. Here you'll find hardware, firmware, tools, tips, and tutorials for modding your rig's performance and appearance. Send us your own mod-related tips and ideas at modding@cpumag.com.

Modding enthusiasts have a penchant for the latest toys. Fast processors, powerful video cards, silent SFF enclosures, and radically lit motherboards are all fair game when it comes to a modder's creative mind. Recognizing the appeal of unconventional customization, an entire industry has emerged to support the community.

Mods & Ends

SOYO Raptor SY-PS400 Adjustable Power Supply

SOYO, traditionally a motherboard manufacturer, is spreading its wings and appealing to the upper echelon of power users with its adjustable 400W power supply ($75) that comes with an external control unit.

SOYO's new power supply comes with its own control unit.

The power supply is a piece of work, sporting standard ATX and 12V auxiliary power connectors, six large four-pin connectors, one small connector, and a pair of SATA power connectors. It also has headers for thermal management and exterior control. Rather than employ a small, noisy cooling fan, SOYO uses a slower, 120mm fan, reducing noise output significantly.

Anticipating heavy draw on the 3.3 and 5V rails, SOYO configured its control unit to monitor and display the status of those two voltages. If your 5V rail dips, you can adjust the Raptor to augment output, compensating for the drop. The control unit also protects against voltages that exceed the Raptor's preset threshold. Further, SOYO includes black and beige faceplates for aesthetic appeal.

We didn't experience any problems with the unit, but our only concern is its +12V rail, which only offers 16A of output. A heavily overclocked processor in a system with multiple drives could have trouble, and there's no monitor to diagnose that particular rail. For comparison, Enermax's EG365P offers 26A on the +12V rail.

CardKeeper Retention Mechanisms

The latest generation of graphics cards are apt at 3D gaming but are also large, heavy, and likely to shift while in transit. The CardKeeper ($3), a deceptively simple looking device built of stiff plastic, helps prevent add-in cards from dislodging by securing a card's front and back.

CardKeeper is available in several versions to accommodate numerous add-in cards. The largest, developed with NVIDIA, is sturdy enough to keep a GeForce FX 5900 in place during transit. You can equip smaller variants with extension arms to keep even the smallest PCI cards locked down. CardKeeper might be an unnecessary for some, but it's an invaluable security measure for LAN gamers.

SilverStone SST-TJ01 Temjin

If you want an affordable case but still value appearance, check out SilverStone's SST-TJ01 ($75); it's a midtower in the form of a 21st century skyscraper. SilverStone boasts about its aluminum front panel, but the chassis is actually steel, making it a 17-pound alternative to some popular all-aluminum cases.

The case features a classy, magnetic hinged door that reveals six 5.25-inch bays and a pair of 3.5-inch bays. It also has a pair of internal 3.5-inch bays for hard drives and many front-panel connectors hidden along the case's side bezel. The SST-TJ01 is available in silver with a black stripe or black with a blue stripe.

The SST-TJ01 is artful and functional.

Fashionably Fresh Firmware

Apple Wireless Keyboard & Mouse v.1.1

Apple says this firmware improves performance and reliability of its wireless keyboard and mouse products.

www.apple.com/support/downloads/wirelesskeyboardmouse firmware update.html

Creative Labs Rhomba Portable Audio Player v.1.05.01

Creative fixed a file system failure issue that could cause the player's memory to be erased in certain instances.

us.creative.com//support/downloads/download.asp

Plextor PX-708A DVD±R/RW v1.06

This firmware rectifies an issue with v.1.05 and improves write performance for DVD and CD-R media.

www.plextor.com/english/support/downloads/firm_708A.htm

Pioneer DVR-107D 8X DVD±R/RW v1.12

This firmware adds compatibility for new 8X media, improves playability on CD and DVD-RAM media, and augments DVD+R and –R write performance.

www.pioneeraus.com.au/computer/firmware.html

by Chris Angelini

The ArtMac

The Tale Of A Mac Modder

My wife, Leontine, will use the ArtMac's sides as a rotating display case for her artwork.

Like an old painting, my Mac was fading. That the 450MHz Sawtooth G4 tower I bought in 1999 was having trouble handling the latest version of Photoshop CS wasn't surprising; the surprising thing was that it ran Photoshop CS at all.

Conventional wisdom says that Mac owners don't upgrade, and they *certainly* don't mod. They go out and buy the latest $1,900 Apple cheese grater, stacked with nine fans and a processor that Steve Jobs probably thinks could bring peace to Iraq if only given the chance.

But just as you don't rip down a fine fresco, I didn't want to trash my artist's machine, which I have dubbed the ArtMac. I wanted it to live up to its true potential—getting a full refit from brush to easel, thus, turning it into a true Photoshop monster. On the way, I'd give it a bit of a mod, too, because a machine that creates art deserves to look the part.

ArtMac started out with these specs:

- 450MHz Motorola PowerPC G4
- Apple Sawtooth motherboard with 100MHz system bus, three PCI slots, and one AGP 2X slot
- 640MB of PC100 SDRAM
- 10GB 7,200rpm Quantum hard drive on an ATA/66 bus
- 40GB external 7,200rpm FireWire hard drive
- ATI Rage 128 graphics card with 16MB of VRAM
- Apple AirPort 802.11b card
- Single-button Apple mouse and keyboard
- A slew of USB peripherals connected via an ugly hub

Clearly, I had some work to do. Could I give this Gothic-era G4 a kick up to the 21st century? Could I do it for less than the price of a comparable new machine? And how much more would it cost to upgrade than a PC in a similar sorry state? To help answer these questions, I recruited Larry O'Connor, president of Other World Computing (www.macsales.com). The Mac upgrade world is pretty small, and OWC is known for producing and selling high-quality upgrade parts at reasonable prices. OWC isn't the cheapest outlet, but its service, support, and quality assurance are impeccable. And hey, this is a Mac. I was willing to spend a little more for a no-worry experience.

To benchmark my progress, I used Xbench (www.xbench.com), a comprehensive benchmarking tool that accurately tests everything from UI responsiveness to disk access speed (but not 3D video performance). I also set up a grueling batch process in Photoshop CS, churning through filters and transforms until the Sawtooth begged for mercy.

Money spent so far: $0; Xbench score: 64.25; Photoshop batch time: 2:30 (minutes:seconds)

Step 1: The Processor

With so many parts to upgrade, I agonized for days about what to throw in the box first. Eventually, I decided on the

The ArtMac's side looks like it would make a great frame for artwork, if only we could get rid of that gray paint.

processor upgrade. The ArtMac's 450MHz G4 chip, while state-of-the-art in 1999, is miserably slow now. Even the system's otherwise slow ATI Rage 128 graphics card is constrained by processor performance. It's a testament to Apple's engineering that OS X 10.3.2 runs on a machine this slow.

OWC's $439 Mercury G4 upgrade looks like a hovercraft, topped by its own heatsink and fan. That is a major differentiator; most other upgrades use the old processor's heatsink, and I wanted a new sink for a new generation. To save money, I picked a Mercury that uses a 1.25GHz Motorola G4 chip overclocked to 1.4GHz, but it is rated for the overclocking, and OWC guarantees it to work at 1.4GHz without unpleasantness.

Flipping open the G4 by its handle revealed a well-organized motherboard crowned by a positively architectural aluminum heatsink. Using needlenose pliers, I wiggled the metal clips off the heatsink and pulled it from the processor daughtercard, a skinny little thing in comparison. I removed three Philips screws, and the daughtercard came off with a neat snap. The OWC upgrade card screwed into its place, and I was done in about 14 minutes.

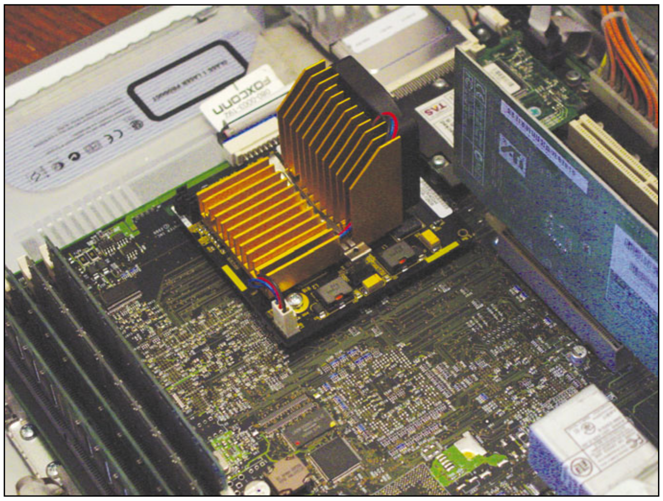

The ArtMac's new processor upgrade, with its own heatsink and fan, snapped into the same socket.

To prepare to mod the case, I had to totally disassemble my Mac.

The computer booted up clean and easy, and the improvement was immediately, startlingly clear. The OS X dock moved more smoothly, Web pages loaded more quickly, and all the benchmarks except drive performance at least doubled. Even the Rage 128 had some life in it.

Money spent so far: $439; Xbench score: 124.76; Photoshop batch time: 1:02

Step 2: RAM, Baby!

RAM is cheap. RAM is fun. More RAM should go to everyone. With four slots, the G4 can handle up to 2GB of RAM but only PC100, thanks to the glacial 100MHz system bus. (Upgrading the motherboard to boost RAM performance would have cost hundreds and only pushed the bus to 167MHz. Not worth it.)

RAM always makes a difference in Photoshop, though, so I decided to add another gig, creating possibly the world's only machine with 1.62GB of RAM. I then tested the machine with both the new processor and the old one. You can use any standard RAM in a Mac, but OWC's own brand RAM is well-known for reliability. I picked up two 512MB chips for $105 each.

Dropping RAM into a Mac is effortless. Pop the case open, admire the swing-out motherboard again, and slide the DIMM into a slot with latches that automatically close when you have it positioned right. Predictably, the RAM did not make any difference to the benchmarks. But it did yank down my Photoshop transform times by about 10 to 15 seconds, with or without the new processor. If my pennies were limited, I would still have saved up for the processor upgrade. On a machine this old, bumping the processor delivered the most bang for my buck.

Money spent so far: $549; Xbench: 123; Photoshop batch time: 0:53

Step 3: The Video Card

Yes, the ATI Rage 128 video card is ancient and pathetic. Guess what? So is the whole Mac video card market. Mac towers can mostly use commodity parts. Memory, hard drives, SATA cards—*mi* interface *es su* interface. But Mac video cards have different ROMs and occasionally different output ports than PC video cards. At press time, only two modern, aftermarket video

Comparing The Macs

Thousands of Mac owners have posted their benchmark results at Xbench.com. Here's how 10 representative models checked out:

Model	Processor	Year Introduced	Average Xbench Score
High-End G5	Dual 2GHz G5	2003	193
Dual-Processor G4	Dual 1.25GHz G4	2003	151
Low-End G5	1.6GHz G5	2003	137
Upgraded ArtMac	1.4GHz G4	2004	136
17-inch iMac	1.25GHz G4	2003	118
$799 eMac	1.25GHz G4	2004	118
Flat-Panel iMac	800MHz G4	2002	73
PowerBook G4 Ti	667MHz G4	2001	67
Power Mac Cube	500MHz G4	2000	66
Original ArtMac	450MHz G4	1999	64
Original iMac	233MHz G3	1998	15

(Source: Xbench.com; lowendmac.com)

cards were available for the Mac: the ATI Radeon 9000 and the ATI Radeon 9800. The 9800 costs $375. (The PC version is $200.)

NVIDIA? I pity ya. Sapphire? Shattered. PowerColor? Powerless. You can count all the Mac video card options on a Muppet's fingers. The ATI Radeon 9800 is a far better card than the ArtMac can handle with its wimpy 2X AGP slot, so I was at a real loss for options.

But wait. It gets worse. The Mac version of the 9000 is loaded down with the useless ADC (Apple Display Connector), which only Apple-branded flat-panel displays use. I needed an S-Video port for the LCD screen mod I had planned, so I got a discontinued Radeon 8500 instead, a dual-head card which sports VGA, DVI, and S-Video connectors. It offers pretty much the same performance as the 9000. (I'm using the 9000's $160 retail price in my calculations, however.)

I did have one other option: Some enterprising Mac hackers have reflashed the ROMs on 64MB PC Radeon 8500 cards to work on Macs (www.coffeehaus.com/radeon8500.html). These run for $50 to $75 on eBay—a big savings. But I'd heard about intermittent problems with deep sleep and dual-monitor support, so I decided to spring for the real deal.

Upgrading to the 8500 didn't make any difference to my Photoshop performance or Xbench results. The old Rage 128 already supports Quark Extreme and Apple's graphics acceleration technology, and the machine won't be playing games that would support the 8500's shaders and transparency functions. All I got out of my upgrade was the ability to span two monitors, which I'd need for my LCD mod.

Kids, don't look—it's a nude Mac.

Money spent so far: $709; Xbench score: 123; Photoshop batch time: 0:53

Step 4: The Hard Drives

A good artist makes the best possible use of materials; a painting on canvas is different from one on paper. When I saw the hard drive bays in the G4, I knew they were a resource I couldn't waste. The Power Mac G4 case is loaded with hard drive bays. The one internal hard drive comes in a bracket that welcomes a second drive on top. There are two other bays, just like that. Sure, there's only one free plug on the IDE controller, but that ATA/66 controller needed upgraded anyway.

First things first: To see how much of a role drive speed played in my machine's sad performance, I got a $60 ACARD PCI Serial ATA controller from OWC and a $90 80GB IBM/Hitachi 7,200rpm Deskstar SATA drive. For Photoshop, the machine doesn't really need the space (the existing 40GB FireWire drive is plenty), but it could use the speed.

The SATA drive snapped neatly into the machine's second bay, but that's when I got frisky. Why use an external FireWire drive, anyway? Sure, I'd get a little drop in performance by switching it to the internal ATA/66 bus, but I had the SATA drive for speed, and those drive bays were begging to be filled.

I backed up the contents of the FireWire drive to the SATA drive and cracked open the FireWire drive's enclosure—two screws and a sliding panel—to reveal a 40GB 7,200rpm Deskstar drive. It would have looked handsome in the third drive bay, but the internal IDE cable was too short. I had to stack the second drive on top of the first. (The bracket was slightly too wide for the second drive, requiring a spacer or two very long screws.) I rejumpered the drive to slave, screwed it into the existing bracket, and, voila, a three-drive Mac.

Finally, I switched Photoshop's swap drive and all my data files over to the SATA drive. I would probably get even better speed with a clean reinstall of Panther onto the SATA drive,

Soak the side panels in alcohol for a while, and the paint begins to dissolve.

but my Panther install CD seems to have gone missing.

Money spent so far: $859; Xbench score: 135.64; Photoshop batch time: 0:51

Step 5: Gravy (Optical Drives, Etc.)

Artists use a lot of ports. My ArtMac started out with only two USB 1.1 ports and two FireWire ports for a printer, scanner, drawing tablet, digital camera, video camera, and iPod Mini. Sure, I could hook them all up with hubs, but installing additional ports using a PCI card is cheaper and more elegant. A SmartDisk PCI card added three USB 2.0 ports and two FireWire ports to the machine for only $35, eliminating the need for hubs. (I know FireWire 800 is out there, but I don't plan on editing any high-def TV footage with this machine any time soon.)

The one-button Apple mouse had to go. I replaced it with Gyration's excellent Ultra GT Cordless Optical Mouse ($75), which offers a right button, a much more elegant look than other multibutton mice, and an incredibly cool gyroscope mode. Holding the mouse like a presentation remote, I can actually use it while waving it around in the air.

Replacing the existing CD-RW with a DVD±RW wouldn't improve performance, but it would let Apple's iDVD software produce beautiful DVDs of artwork. iDVD only supports internally mounted Pioneer drives, so I got the latest and greatest Pioneer DVR-107 ($129), which writes to both DVD+RW and DVD-RW discs at 8X.

Installing the drive was, as with everything else, a breeze. The plastic faceplate covering the optical drive snaps off by pressing two tabs inside the case, revealing two screws. Undo the screws, and a sled containing the optical drive and Zip drive neatly slides out (covered by a pointless metal endcap, which I threw away.)

The existence of the Zip drive created a moment's doubt. I'd love to have two optical drives in this machine, and I'd seen a mod to replace the Zip with another optical drive (www.shanescot.com/photos/G4_mod), but that mod requires serious metalworking to turn the smaller Zip drive sled into a larger optical drive case. Oh, well. Slide the CD-RW out of the sled, drop the DVD±RW in, reattach the ATA and power cables, and I'm good to go.

Money spent so far: $1,098. No effect on performance.

Interlude: How Am I Doing?

By this point, I was wondering whether I'd saved any money vs. buying a new Mac. So I dipped onto Xbench's system-comparison Web site to see how other systems compared. Apple no longer sells G4 towers. For $999, I could get an all-in-one eMac with a 1.25GHz G4, ATI Radeon 9200 with 32MB of VRAM, 80GB of ATA/133 storage, 8X DVD-RW drive, and 256MB of RAM that would benchmark in the 120 range. I've got 10% better performance for $100 more than the base eMac, but I also have more storage, more RAM, more ports, a fancy mouse, wireless networking, and a dual-head video card with more VRAM.

For $1,799, meanwhile, I could get a hulking G5 running at 1.6GHz with an NVIDIA GeForce FX 5200 Ultra, 80GB of SATA storage, 4X DVD-RW drive, 256MB of DDR333 RAM, and an 800MHz frontside bus. It would benchmark in the 140s, for a slight gain in performance over the ArtMac, but I've just saved $700.

Don't even get me started on how these upgrades compare to similar purchases for a PC. The RAM, SATA card, hard drive, and USB card I used are PC-standard equipment. But for a PC, I could have grabbed an NEC ND-2500A 8X DVD writer for a mere $87 from NewEgg.com, scored a PowerColor Radeon 9200 SE for $47, and grabbed a 2.8GHz Intel Pentium 4 for $150. I'd save a total of $444 on these upgrades if I were doing them on a PC. That's the price you pay for subscribing to the Cult of Steve.

The Case Mod

Remember those Muppet fingers? You can count Mac case mods on those, too. Mac owners usually fetish-ize their machines; they don't cut into them. A few brave souls have, though. (See the "Better Men Than I" sidebar for some of the best mods out there.) This was my first mod (hey, I'm a Mac user), so I decided to go easy. I would bleach the

ArtMac's gray side panels so I could use them as gigantic frames for artwork and stick an LCD in the machine's top to display video art.

The Mac case is a masterpiece of obsessive over-engineering. Screws aren't good enough for Steve Jobs; he needs to have screws, tabs, *and* flanges, all together wherever possible. One gray side panel, for instance, is attached to the case with four Allen screws but also with plastic tabs that tuck in under the fan. So, you have to remove the fan to get the panel off. While removing the fan, I realized it was screwed into the case like with any normal computer but also tucked in with metal tabs so it snaps into the power supply.

Once the computer was completely disassembled, I turned to installing the LCD. I picked up a 4-inch NTSC LCD on eBay for $60; it came with a mini jack output carrying both composite video and mono audio and its own power supply. I tested it by plugging it into a PC with an ATI Radeon 9000 video card, into a DVD player, and into a cable box. It seemed everything worked.

Using a cordless Dremel, I cut identical holes in the plastic and metal panels on the top of the case. The LCD had to be a little off-center, as a tempered steel security strut (meant to let you chain your Mac to a desk) runs the entire length of the case. (More over-engineering.) The Li-Ion cordless Dremel is powerful, but I had to recharge it twice during the cutting process.

With a Philips screwdriver, I stripped the LCD out of its plastic bezel, and it fell apart into three pieces: the LCD, the backlight, and a control board that was unfortunately glued to the plastic back of the case. I used epoxy to glue the backlight to the LCD and stuck the control board back on with ordinary packing tape. Then, very nervously, I epoxied the whole unit face-up inside the case. Between the metal and plastic layers of the case, I put a clear, circular disc of plastic to protect the new screen.

My cordless Dremel dug into the top of the metal case to make room for the LCD.

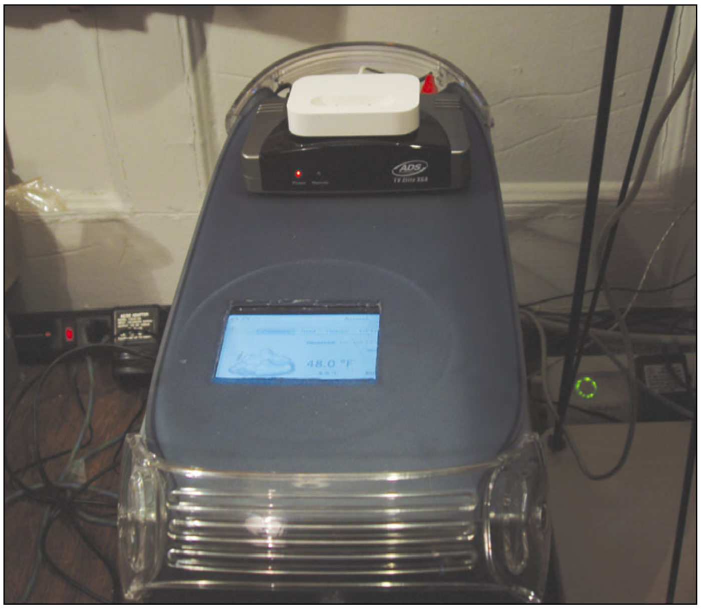

The ADS scan converter makes a great little stand for my iPod Mini dock. In the mornings, the LCD shows the weather.

I left a mess, but this was my first mod. The LCD's power and video cables ran out an empty PCI slot, where I hooked the power adapter up to an external switch. If I had more hardware experience, I could have cut away the plastic from the control board, plumbed the LCD's power supply into the Mac, removed the mini jack, rewired the output to use an RCA plug, and drilled holes into the case to fasten the assembly with screws rather than glue.

The only real downside of my approach (other than internal elegance) is that turning off the Mac doesn't power down the LCD. So I plugged the LCD, Mac, and the scan converter I ended up needing later into a power strip to use as a main switch.

Soaking the G4's side panels in isopropyl alcohol for a few hours stripped the gray paint, leaving translucent plastic panels with a blue plastic apple in the middle. (More over-engineering.) You can repaint the side panels with model-car enamel or using Krylon Fusion spray paint, but I decided to use it as a display case for my wife's artwork instead. I snapped off the center tabs hooking the side panels into the case, leaving the panels attached by four screws, so I can replace the artwork frequently.

The Unforeseen Horror Of ATI's Mac Drivers

Everything looked to be going great. With my new machine ready to go, I plugged the LCD into the video card, and—nothing. Huh? I plugged the video card into another TV, and it worked. I plugged the LCD into a VCR, and it worked. I plugged the card into a TV, booted the Mac, pulled the cord from the TV, and plugged it into the LCD. It worked. The computer could drive the LCD, but it just refused to boot with the LCD. I finally called ATI and discovered the miserable horror of ATI's Mac drivers.

It's not enough that there are only two aftermarket video cards for the Mac. The drivers are, to be blunt, lousy. The Mac Radeon's video-out port only turns on if it detects a 75 ohm circuit connected, and according to ATI, the mini jack input on my LCD can't close a 75 ohm circuit. All PC Radeon drivers have a simple toggle switch that forces the TV-out on. Not so with the Mac drivers. And ATI has no plans to add such a switch. I spent a week puttering with various cables and even a Radio Shack signal amplifier. Nothing worked, and ATI tech support was unhelpful. So I buckled and ordered a TV Elite XGA scan converter ($129) from ADS Technologies.

ADS' scan converter is supposedly for presentations; it pipes a VGA output into a large-screen TV. But it also works with a very small-screen LCD. ADS' tech support bent over backward to help me, a refreshing contrast to ATI's modder-unfriendly tech guys. As a bonus, my ArtMac can now output video to a VCR. The LCD looks a little fuzzy, but I'm not about to do word processing on it. Instead, I'm writing an AppleScript that will rotate various pieces of artwork and interesting information onto the screen: video art in the afternoon; weather information in the morning; and spooky, swirling colors at night.

Making Art History

The ArtMac is a unique art-ifact, one of very few nonstandard Macs in the world. And it runs Photoshop CS with (relatively) blazing speed. So yes, you can upgrade a Mac. Yes, it's fun. And no, it isn't cheap. But it's considerably cheaper than buying a new Mac.

With a very old machine, the processor was my No. 1 bottleneck, choking everything else, even an ancient video card and ATA/66 IDE bus. Boosting RAM provided excellent bang for the buck in Photoshop, though not in applications that dealt with smaller files. Upgrading my hard drive bus, on the other hand, provided a surprising across-the-board boost.

The Mac video card market is stuck in the dark ages. There are only two aftermarket cards out there, they're stunningly expensive, and they provide little benefit to nongamers other than a buggy dual-head capability. Unless I needed to run two monitors or Unreal Tournament, I'd save my money on that upgrade.

The big surprise wasn't that I could mod a Mac. It's that so few people are doing it. Mac users are supposed to be creative, imaginative types. Hopefully, the ArtMac will inspire a few to turn their own Macs into unique works of art.

by Sascha Segan

Better Men Than I: The Real Mac Modders

Mac modding probably started in 1993 with Andy Ihnatko's Macquarium. Ihnatko's original plan turned a classic 512K Mac into a functioning aquarium. Now there's even a business, MacAquarium (store.macaquarium.com/aquaria/customer/home.php) devoted to selling kits and finished aquariums based on original Macs or the 1998-era iMac.

In terms of actual functioning mods, though, my work pales before that of Kent Salas, probably the greatest Mac modder alive today. His Project Pyramac (www.riscx.com/pyramac) started with a machine similar to mine—a G4 upgraded with an OWC 1.4GHz daughtercard, but it went much further. Salas ditched the original case, building a translucent, pyramidal acrylic case instead. He painted various internal components with ultraviolet-sensitive paint and stacked ultraviolet tube lamps inside the case. The result looks nothing like a Mac; it radiates an eerie light and displays its innards like a jellyfish. It's brilliant.

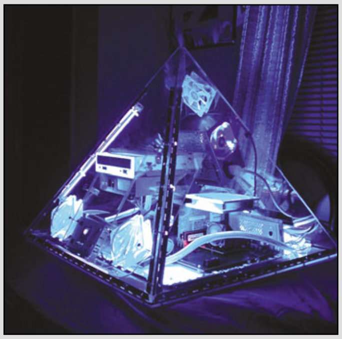

Kent Salas' Pyramac used UV-sensitive paint for a unique glow.

Kirk Newcombe's Xyris Project (www.xyris.vindicated.org) is a simpler mod but marked by attention to detail. He started with a G4 and removed the paint from the case. Newcombe then spent considerable time adding yellow paint highlights and blue LEDs. He repainted the speaker, power button, and DVD tray. His only major performance mod was replacing the G4's internal Zip drive with a 10GB hard drive.

Several modders have attacked the G4 Cube. Check out CubeOwner (www.cubeowner.com). Cubes with wooden interiors, chromed Apple logos, and appearances like R2-D2 are among the most prominent mods. The only aftermarket Mac case is for the Cube. The PowerLogix PowerCube ($139; www.powerlogix.com) expands the size of the Cube and tops it with a G5-like aluminum grill for modders who want to pack in processor or video card upgrades.

Modders who like the look of Macs but want to run Windows have gone another direction, stuffing PC innards into old Apple cases. We covered one such mod in our October 2002 article "One Rotten Apple," where we fit a gaming PC into a G4 case. Attila Juhas' SE3000 (http://users.adelphia.net/~attilaerika) takes the concept to the next level. He fit a compact gaming PC in a mid-1980s Mac SE case. Thanks to the tiny MicroATX motherboard, Juhas managed to pack a 3GHz P4 machine with 1GB of RAM and an ATI Radeon 9700 PRO video card into the cute little box. He cut a DVD drive door into the base and sandblasted away most of the internal monitor, which now functions as a case window. Professional-quality logos and lights complete a truly beautiful mod. ▲

Take the ViewSonic® screen test!

THE TABLET PC V1100. MORE POWER. MORE FREEDOM. MORE STYLE.

There's a revolution in PC portability that's as important to your freedom and productivity as the cell phone or laptop. The convergence of the PC with high-resolution display technology is realized with power and mobility in the ViewSonic Tablet PC V1100. If your home is often your office and the road is your conference room, the Tablet PC V1100 is the perfect combination of office manager, legal pad, portable desktop, and productivity networking tool. It offers all the functionality of a Windows® Office™ environment in an unobtrusive, ultra-stylish 3.4 pound profile. With a big write-on screen and superior handwriting recognition, the V1100 does much more than a laptop. Take handwritten meeting notes, or sketch big ideas and immediately share them with your team via integrated LAN and Wi-Fi. Any application that runs on a laptop will run on your Tablet PC. Combined with a ViewSonic projector like the new 2.2 pound PJ250 micro-portable, you've got an incredibly powerful mobile presentation solution under six pounds. Take a screen test today at your ViewSonic reseller or visit **ViewSonic.com** and see why the Tablet PC V1100 is a runaway hit.

Add our compact projector for mobile presentations, or with keyboard included and optional dock, the V1100 easily switches from portable to desktop duty.

Freestyle

Specifications and pricing subject to change without notice. All logos, trademarks and registrations are the sole properties of the their respective companies and are used by permission.
Copyright © 2003 ViewSonic Corporation. All rights reserved. [11042-00B 03/03]

ViewSonic®
See the difference.®

Mad Reader Mod

The Fragmatatron: A Rusted-Out Beauty

In case you haven't noticed, custom paint jobs are in. Sure, there is more than one beige box holdout, but manufacturers regularly parade flashy rigs through our doors. We like shiny objects as much as you, but Derek Thanstrom, aka GK>Clockwerk Rat, showed us that you don't need auto paint or an army of LEDs to catch our attention. In fact, a dab of rust-colored paint here and there can turn that beige box into a work of art.

Well, maybe there is a little more to the Fragmatatron than paint. Inspired by his father's love of WWI electronics, Thanstrom used rivets, wood screw plugs, gauges, and a custom paint mix to give his computer its gritty "been around longer than you but still runnin' fine" look. He also mounted a plasma ball in the huge Antec SX1240's upper drive bays, hardwiring the 12V device to a custom bus bar. The switches control the Fragmatatron's gauges, cold-cathode lights, and plasma ball.

Thanstrom spent a month scavenging for the Fragmatatron's components and another two-and-a-half weeks bringing it to life. The system includes an AMD AthlonXP 3200+, ATI Radeon 9800XT, and 1GB of Corsair TWINX 1024-3200LLPRO DDR RAM. The hard drive fans originally cooled two ultra-hot 10,000rpm SCSI hard drives; Thanstrom recently replaced those drives with a 120GB SATA Seagate Barracuda. The Fragmatatron won its DFI LanParty nForce2 Ultra 400 motherboard in a LANParty.com modding contest. Thanstrom's plans for future Fragmatatron mods include a belt-driven brass fan and a skeleton key power switch. ▲

by Joshua Gulick

Want your own rusty paint job? Subscribers can see instructions at www.cpumag.com/cpujun06/rustbucket.

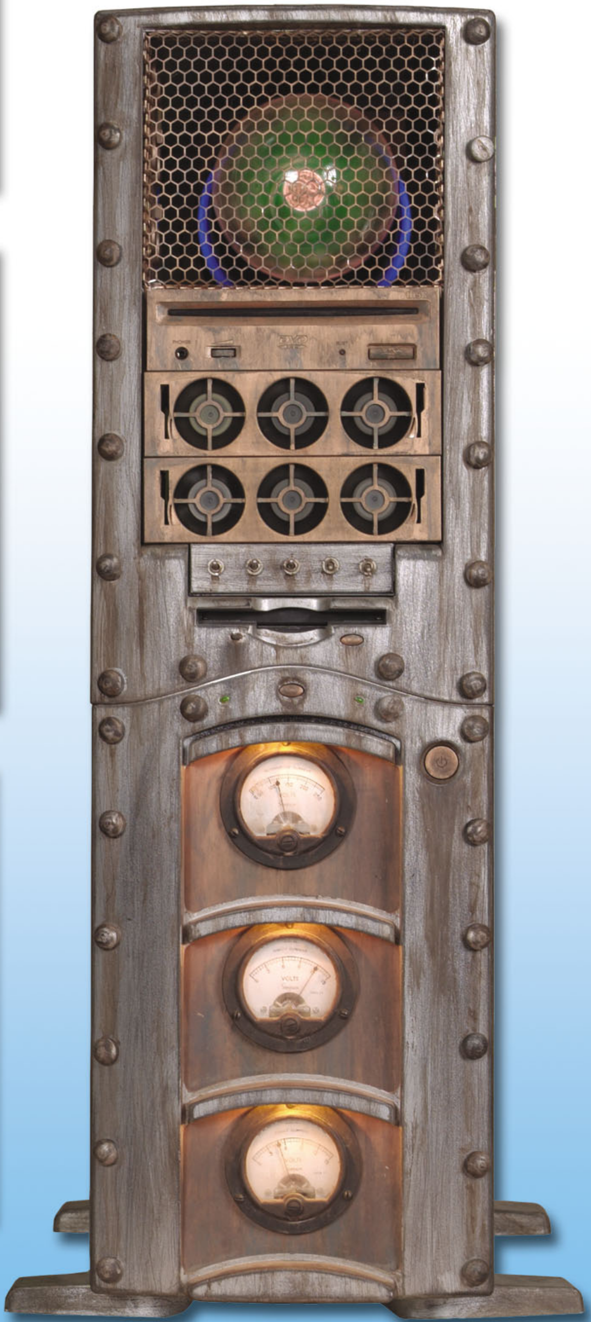

Advanced Q&A Corner

Get informed answers to your advanced technical questions from *CPU*. Send your questions along with a phone and/or fax number, so we can call you if necessary, to q&a@cpumag.com. Please include all pertinent system information.

Each month we dig deep into the mailbag here at CPU *in an effort to answer your most pressing technical questions. Want some advice on your next purchase or upgrade? Have a ghost in your machine? Are BSODs making your life miserable?* CPU*'s "Advanced Q&A Corner" is here for you.*

John Bjacek asked: **I currently have an Athlon XP 2200+ with an ECS K7VTA3. The motherboard supports a maximum FSB of 333MHz. I would like to upgrade the CPU in my computer to the best one possible and was wondering if I would get more performance from an Athlon XP 3000+ with a 333MHz FSB or by running an Athlon XP 3200+ with a 333MHz FSB instead of its default 400MHz?**

A: The short answer to your question is that the Athlon XP 3000+ will give you the best performance with your motherboard, but there are quite a few other factors you should consider. If you run the Athlon XP 3200+ with a frontside bus of 333MHz, you'll be dropping its operating frequency down to 1.82GHz, which is the same clock speed as an Athlon XP 2500+. The 3200+ is normally clocked at 2.2GHz (11 x 200MHz FSB = 2200MHz), but by decreasing the FSB frequency, its clock speed is lowered. The Athlon XP 3000+, on the other hand, has a higher multiplier, so when it's installed in your motherboard, it will operate at its proper default clock speed of 2.16GHz (13 x 166MHz).

At about $160, however, we think your upgrade may not be worthwhile. With an Athlon XP 3000+ based on AMD's Barton core, you'll get yourself an additional 256KB of L2 cache and a 16.5% boost in clock speed, but the overall effect on real-world performance may not be perceptible. Your benchmark scores would definitely increase, but if you can't "feel" the increased performance in your day-to-day computing, the $160 expense isn't practical, in our opinion. You'd be better served by overclocking your existing CPU or by saving a few more dollars and upgrading your motherboard and processor to, say, an Athlon 64 or midrange Pentium 4.

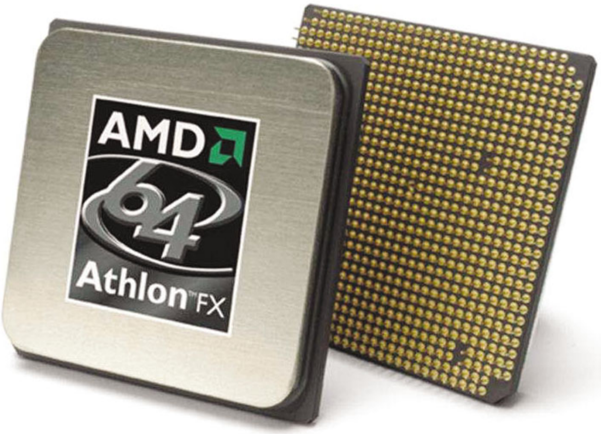

The AMD Athlon 64 FX-53 is arguably the fastest desktop processor available today, but multithreaded apps might do better with the Hyper-Threading architecture of a P4.

If a motherboard upgrade is out of the question and you're dead set on upgrading your existing rig, look into a mobile Athlon XP. AMD's mobile Athlon XPs are basically handpicked processors that run at the same frequencies as their nonmobile counterparts, but with lower default voltages. These mobile processors have proven to be very capable overclockers, and you can find them for less than $100 at a few online retailers. For example, the mobile Athlon XP 2500+, which you can find for about $90, would almost certainly overclock to the 2 to 2.2GHz range, giving you all of the performance of a standard Athlon XP 3000+ at a significantly lower cost.

Jeff Woolson asked: **My boss has tasked me with acquiring a PC for the company to be used to process mathematical calculations. The machine we currently use is an Intel P4 2.2GHz, and the processing takes four to eight hours for each work unit. What I need to know is which line of processors is better suited to number crunching, Intel P4 or AMD Athlon 64? This machine will not be used for gaming, just processing RF antenna modeling, subspace variables, and RF radiation dispersion patterns.**

A: This is an excellent question, Jeff, and there have been plenty of analysis and comparison benchmarking on the topic around the 'Net and in print as of late. (Check out pages 21 and 22 of this issue.) The SPEC Viewperf tests in this article may be more closely akin to the type of work you are running on your machines. Both processor architectures have their merits depending on the applications and software you're running most.

Which line of processors is better suited to number crunching, Intel P4 or AMD Athlon 64?

From a raw computer performance standpoint, if money is no object, AMD's Athlon 64 FX-53 (or FX-51, for that matter) seems to hold the current top-end performance lead over the P4. However, if the applications you are running either make heavy use of SSE instructions or are multithreaded apps, Intel's P4 could provide even better performance. The Pentium 4's Hyper-Threading capability, assuming you are running Windows XP Professional, allows it to perform like a multiprocessor machine but with one CPU because it allows the processor to handle two threads to distribute the workload in a parallel processing approach.

The P4's SSE instructions are specifically targeted at floating-point calculations to accelerate things such as multimedia, gaming, digital video, and audio processing. If the type of math that is done in your modeling is FP intensive and is optimized for SSE, SSE2, or Intel's latest SSE3 instruction set, the P4 will most likely be a better choice. The Athlon 64 also supports SSE and SSE2, but, in general, its performance with SSE-enabled code isn't as robust as the P4's.

Finally, if the applications you run are larger programs that like to consume memory resources, choosing one of the new Athlon 64 FXs with their onboard memory controller and large 1MB L2 caches may help a bit. In this case, the Pentium 4 Extreme Edition (with 2MB L3 cache) or possibly the P4 Prescott CPU (with 1MB L2 cache) may also be good choices.

Brenner Fox asked: The computer in question is an Athlon XP 2200+ machine with 512MB DDR running Windows XP Pro. It was built for a friend's business, so it's a mission-critical unit. It has proven to not be very reliable, however. The problem is that the system will work fine for a number of days, then suddenly it will no longer boot into Windows. And if we try to go into safe mode or load the last known good configuration, the monitor turns black, and the system will restart.

The problem seems to be unrelated to heat, RAM, or a CPU malfunction because a reinstall or ASR (Automatic System Recovery) instantly repairs the situation. As a second option to using ASR, I decided to make a restore image using Drive Image. I received numerous errors trying to create an image, though, so I ran Check Disk and repaired the errors. Then when I created an image, I didn't receive any more error messages. Unfortunately, the problem reoccurred the very next day, and when I tried to restore the image, I got a lot of error messages relating to lost files and bad clusters.

This latest discovery made me think that something may be wrong with the hard drive, but there are no strange sounds coming from it at all. Except for the errors mentioned earlier, it seems like a properly functioning drive. In a new twist to the story, I just repaired another older computer networked to this one that had the identical problem right before this one crashed again.

A: Brenner, my friend, this sounds like a really bad day at the office. We feel your pain. Drive corruption, which seems to be the cause of your crashes and nonbootable failures, is a bit of a traumatic experience for most PC technicians and enthusiasts because it sometimes can take a long time before you determine that data corruption is the root cause of your headaches.

Because you've noted that this identical anomaly has occurred on two machines that are networked, we suspect a virus or worm has found its way into the office network (God knows there are a million variants out there these days) and is propagating itself back and forth between the two machines in question. The only way to rule out this scenario is to disconnect both machines from the network and do a clean format and install of WinXP on both machines. Also, it could be that this virus is corrupting the master boot record of your drives. Therefore, before you do a format and reinstall XP, reboot with your XP CD and select the Repair A Windows XP Installation With The Recovery Console option. Once in the recovery console, type **fixmbr**, which will wipe your master boot record clean. Alternatively, you could boot to a trusty old Windows 98 boot floppy and type **FDISK /MBR** to do the deed. You can then reboot from your WinXP CD, create an NTFS partition with XP, and reinstall. You'll then be sure that you are totally clean, from your Master Boot Record on down to the new partition you've created.

Don't reconnect either machine to the network until they both have fresh installs with clean master boot records. Also, check the hard drives for errors and repair them with Scandisk (XP does this automatically) before you reinstall, obviously. If there are other machines on this network, they too could be propagating a virus, so they may need to be disconnected and scanned for viruses with an updated virus program such as Norton AntiVirus or McAfee Virus Scan. You may also have to perform the same fixmbr, partition, and reinstall procedure on those machines if they are corrupted and your virus scan programs can't clean them.

We Feel Your Pain.

Drive corruption . . . is a bit of a traumatic experience for most PC technicians and enthusiasts because it sometimes can take a long time before you determine that data corruption is the root cause of your headaches.

If this doesn't clear up your data-corruption issue, you may indeed have a couple of drives going bad on the two machines you noted here. That's just plain tough luck, but in the wild world of personal computing, anything can happen.

David Snyder asked: I intend to build a home theater PC soon, but I'm uncertain what to look for in a video card. I'm not the least bit interested in gaming but would like high-performance playback of DVDs and eventually high-definition television. What is important for this type of video playback? Do any of the cards offer anti-aliasing or other image processing that would benefit DVD playback? I plan to use both S-Video and DVI outputs.

A: Virtually all mainstream and high-end video cards on the market today accelerate DVD playback and incorporate some form of adaptive de-interlacing technique to keep feathering and pixelation of video at a bare minimum. Most of them have S-Video and DVI outputs, as well. For your situation, two products immediately come to mind: an ATI Radeon All-In-Wonder or any card based on S3's new DeltaChrome S8 chipset.

We recently tested the S3 DeltaChrome S8 and S8 Nitro and found that they had very good 2D image quality and DVD playback, but their gaming performance was subpar. The DeltaChrome also offers integrated HDTV and Standard TV encoders with 480p/720p/1080i/1080p component (YPbPr) output. What the S3 DeltaChromes don't offer, however, is component input or TV-tuners, which would require the purchase of a separate card should you also be looking for PVR (personal video recorder) functionality.

ATI's All-In-Wonder 9600XT is a great all-around card for gamers and video buffs.

If you're looking to build a home theater PC to act as a PVR, an ATI Radeon All-In-Wonder would probably be the best fit. Our choice would be the relatively new All-In-Wonder Radeon 9600 XT. The AIW 9600 XT comes with a wireless remote and has great DVD playback. It also has composite and S-Video inputs and outputs, along with built-in TV and FM tuners. ATI's Multimedia center software, which works in conjunction with the All-In-Wonder products, would give you almost all of the PVR functionality found in set-top boxes such as a TiVo or ReplayTV. About the only thing "missing" from the AIW 9600 XT is HDTV output, but you can remedy that with an additional card such as ATI's new HDTV Wonder.

Deciding which cooler to choose depends on what is more important to you, noise or overclocking performance.

Nightingale asked: I am in the market for a good CPU cooler and am torn between two particular models: the Thermaltake Extreme Volcano 12 (A1745) and the Zalman CNPS7000A-AlCu. I don't know which would be better in my situation. I will be using it on an ASUS A7N8X-E Deluxe motherboard with an Athlon XP 2800+ and plan to do some overclocking.

Zalman's CNPS7000A-AlCu CPU cooler is designed for near-silent operation.

A: Either one of these coolers would work well, but deciding which one to choose depends on what is more important to you, noise or overclocking performance. We haven't done a side-by-side comparison of these two coolers, but based on their specifications, it seems like the Thermaltake Extreme Volcano 12 is the more capable of the two. The Thermaltake cooler is made completely of copper, which would conduct heat better than the aluminum/copper hybrid Zalman CNPS7000A-AlCu. The Thermaltake cooler is also much heavier at 708 grams vs. the Zalman cooler's 445 grams, and its single, three-pronged clip would make it easier to install. Installing the Zalman cooler would require you to remove your motherboard and install a pair of brackets that help hold it in place. But this more complex installation can also be considered an advantage; once it's installed, it's more secure.

Another advantage of the Zalman CNPS7000A-AlCu is its near silent operation. Both of these coolers have user-adjustable, variable speed fans, which make it easy to configure each one for maximum performance or low-noise operation. However, when dialed up to their respective maximums, the Zalman cooler is significantly quieter at 25dB; the Thermaltake Extreme Volcano 12 is much louder at 48dB. And because the Thermaltake cooler is so much heavier and mounts using a simple clip, we wouldn't advise moving your case once it's installed. Should you frequently transport your PC to LAN parties or other similar events, there is a very real chance of the Thermaltake cooler breaking the tabs on your socket because it is so heavy. The Zalman cooler's custom mounting hardware eliminates the possibility of it breaking the socket while it's in transit. ▲

by Dave Altavilla and Marco Chiappetta, the experts over at HotHardware.com

Cooligy

Making Watercooling Small & Efficient

As engineers and scientists work at developing new methods of cooling the CPU, it's somewhat surprising that they haven't looked more often at borrowing ideas from one of the greatest temperature-regulating machines on earth: the human body. The body, in all but extreme conditions, is able to maintain a steady core temperature by dissipating heat or conserving heat, depending on the prevailing factors. One way the body dissipates heat is by using its complex series of blood vessels and tiny capillaries to carry heat away from the core of the body through the bloodstream to the skin, where the heat dissipates through the skin into the air. The capillaries are effective at transferring heat because they are small and because they can be extremely close to the skin, unlike larger blood vessels.

Cooligy is using a form of this idea for dissipating heat—all without making you or your PC sweat. Cooligy's Active Microchannel Cooling uses liquid inside tiny channels to carry heat away from the core of the computer (the CPU) and dissipate the heat outside the computer body. Cooligy, founded in 2002 in Mountain View, Calif., is hoping its Active Microchannel Cooling technology will eventually replace the heatsink and fan used to cool most CPUs today.

The Need

Current standard cooling technologies for CPUs—which typically include a heatsink and a fan—have worked great for standard processors for several years, but this cooling technique cannot keep up with the ever-increasing heat generated by ever-smaller CPUs that contain millions of miniscule transistors. On the other hand, overclockers have been dealing with heat problems with CPUs for several years and have helped spur the development of watercooling technologies.

The mixture of higher clock speeds and higher density transistors leaves the newest CPUs requiring more power, which generates more heat. As CPU manufacturing now has shrunk to 90nm, the heat problems continue to climb. CPUs that run at temperatures beyond recommendations are subject to computational errors and, if the heat is intense enough, physical damage and failure.

As CPUs generate more and more heat, the traditional heatsink and fan combination isn't practical. With no changes in cooling technology, heatsinks would need to expand and greatly increase their surface area to adequately dissipate the ever-growing heat, which increases cost and limits reliability. Additional fans might be needed, too, adding to the stray noise the computer case generates.

Another problem with today's CPU chips is hot spots on the die, which are tiny, concentrated areas on the chip that run extremely hot as they perform calculations. A cooling system that attempts to treat the CPU as a whole—such as a large heatsink—cannot easily cool these hot spots as they come and go. A watercooling system that uses tiny channels

Active Microchannel Cooling Overview

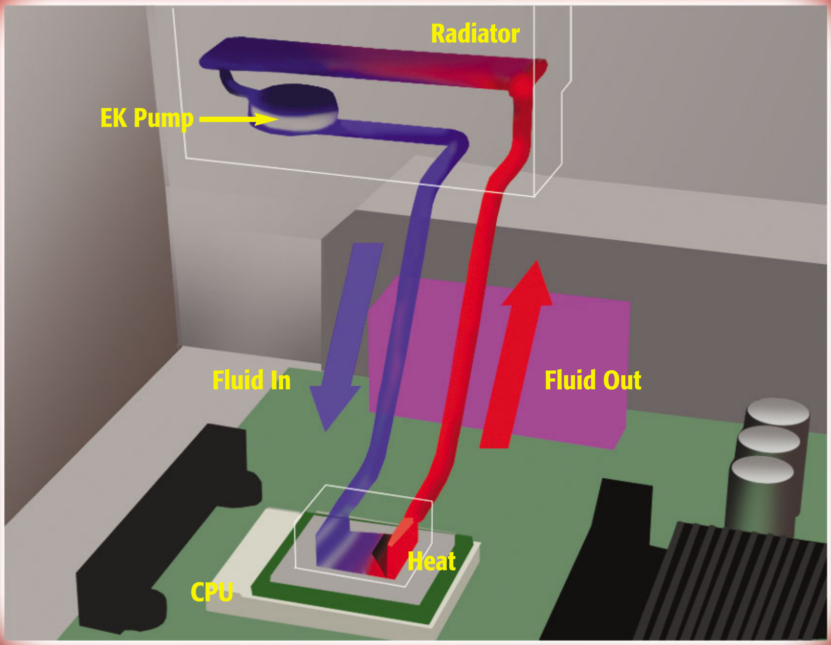

The diagram above shows the flow of the water inside Cooligy's fully enclosed Active Microchannel Cooling system. As the water inside the heat collector absorbs heat from the CPU (lower area) it flows through the system's tube (red tube) to the radiator (upper area). After the fluid releases the heat, it enters the EK (electrokinetic) pump, which returns the now cooled water (blue tube) to the Active Microchannel Cooling chip (or heat collector), where the process begins again. (For more details on the EK pump and on the heat collector, see the "Quiet Cooling" and "Microchannels Magnified" graphics, respectively.)

Source: Cooligy

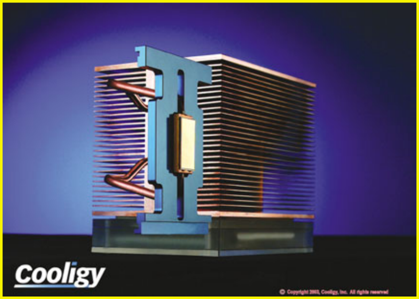

Cooligy's Active Microchannel Cooling technology uses a heat collector (shown here in gold). The heat collector connects to the CPU through the mounting plate (shown in blue). The water in the system moves through the pipes to the radiator (shown in orange), carrying the heat away from the CPU. (Photo courtesy Cooligy)

can more readily access any hot spots on the CPU.

Watercooling technologies successfully deal with many of these problems. Early watercoolers were bulky, however, and some watercooling products weren't reliable enough. Cooligy says Active Microchannel Cooling technology will end the problems while providing the cooling technologies needed to handle today's newest processors and future CPUs, too. Here's a look at how the company's technology works.

The Technology

When developing Active Microchannel Cooling, Cooligy designers took some of their cues from the design of the CPU itself. Active Microchannel Cooling starts as a piece of silicon that contains etchings similar to those on a CPU chip. But instead of carrying electric impulses as they do in a CPU chip, the etchings on the Active Microchannel Cooling silicon actually carry tiny streams of water that cool the adjacent CPU. Cooligy has specially treated the water to prevent growth of any organisms.

As the streams of water flow through the silicon's etchings, they absorb heat energy radiating from the CPU. The streams then carry the heat to a radiator, where a fan or a passive cooling heatsink dissipates the heat into the surrounding air. The Active Microchannel Cooling system is completely self-contained, and the system's fluid doesn't evaporate, making it different from many standard watercooling technologies. (Obviously, the Active Microchannel Cooling system is also different because it operates on a tiny scale compared to most other watercooling technologies.)

Cooligy also uses an EK (electrokinetic) pump with its system. The EK pump changes the electrical charge of the fluid, causing it to move naturally through the system. The EK pump has no moving parts, meaning it creates no noise, and it has an extremely low risk of breaking down.

Cooligy hopes its Active Microchannel Cooling system will eventually replace the traditional fan and heatsink used to cool CPUs.

In the example at top right, the CPU contains no cooling technology, and it would quickly overheat and fail. Heat dissipation (red arrows) is poor.

In the example at center right, the heatsink pulls the heat (yellow arrows) away from the CPU. The fan then blows air across the heat sink, causing the heat to dissipate to the sides. However, the heat remains inside the computer case, causing problems for other components.

In the example at bottom right, the fluid in the Active Microchannel Cooling system efficiently pulls heat away from the CPU and carries it away to the radiator (large red arrow), where it's released outside the computer case. The fluid then returns after it's cooled (large blue arrow). Because the heat collector can sit atop the CPU, its ability to draw heat from the CPU is maximized.

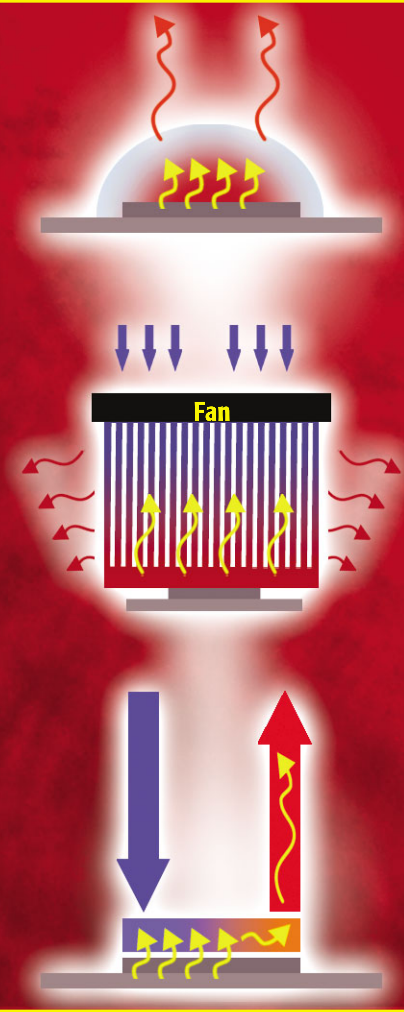

Graphics & Design: Carrie Benes & Jason Codr

Industry Reaction

Active Microchannel Cooling has plenty of industry analysts excited, thanks to its reported cooling potential of 1,000 watts per square centimeter, which should easily meet the needs of current and future processors and address any hot spot problems on the CPU. The reported price of the Active Microchannel Cooling package—about $30 per unit—is also contributing to the excitement.

Cooligy should have the technology available sometime around mid-2004, but it's expected that Active Microchannel Cooling will only be available to OEMs for a while. Cooligy will initially aim Active Microchannel Cooling at workstations and other high-performance computers. It's unclear when or if the cooling technology would be available to individuals and desktop and notebook PCs. However, Cooligy has indicated that it hopes to eventually sell its cooling products to individuals in the enthusiast market and to those with traditional PCs.

by Kyle Schurman

Subscribers can see www.cpumag.com/cpujun04 for more.

Quiet Cooling

The EK (electrokinetic) pump in the Active Microchannel Cooling system provides a quiet method for moving water. The sample system below isn't drawn to scale, and the fan is optional, but the illustration of the EK pump demonstrates the EK effect. With positive voltage applied to the fluid at one end, the hydrogen ions in the water become positively charged. They are then repelled by the negative charge in the glass wall (which is the EK effect). As the positive hydrogen ions repel, they move parallel to the wall and force the water to move close to the glass walls, which propels the water through the system. The water in the EK pump does not go through a phase change, as is common with other watercooling systems.

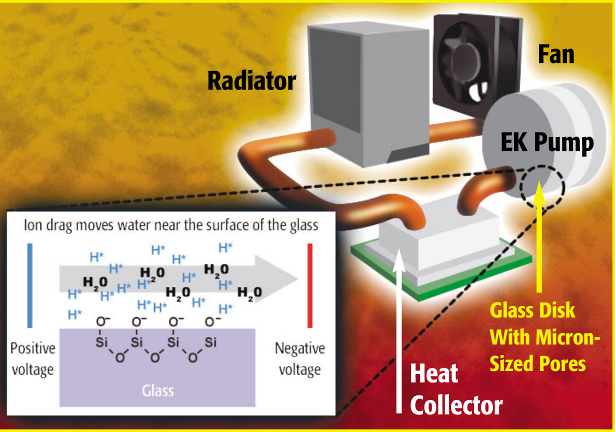

Microchannels Magnified

The microchannels in Cooligy's Active Microchannel Cooling technology are etchings on a silicon chip. Each microchannel houses water that can carry heat away from the CPU. Because the microchannels are so small (usually between 20 and 100 microns in width, depending on the application), they are especially effective at dealing with tiny CPU hot spots. The heat collector works well because it can sit directly atop the CPU; the heat has to move less than 1.5mm before the water begins carrying it away. Several microchannels typically are within a few millimeters of the hot spot, giving it maximum cooling capabilities. When manufacturing the heat collector and the microchannels, a company will predict the locations of hot spots on the CPU and then maximize the cooling efficiency of the heat collector in those areas by making the microchannel walls as thin as possible, which increases the surface area covered by water, thereby maximizing cooling efficiency. Cooligy's research has shown that using microchannels can cool CPU hot spots up to 1,000 watts per square centimeter.

The width and depth of the microchannels is the key to achieving maximum cooling efficiency. As shown at left, when all other factors are identical (such as the microchannel height and depth and the fluid density), an extremely narrow microchannel of about 10 microns can achieve almost a 20 degree Celsius cooling difference versus a microchannel that's 100 microns wide (which is about equal to the diameter of a human hair). The microchannels are about three to 10 times as tall as they are wide to achieve maximum cooling efficiency (as shown in the magnification of a microchannel system at above right). This narrow, tall configuration allows the fluid to heat completely before exiting the microchannel, allowing for maximum efficiency.

The idea of using microchannels to remove heat from a CPU initially appeared in a Stanford University research paper in 1981. Much of Cooligy's initial research occurred at Stanford, and Active Microchannel Cooling actually debuted at Stanford. The Cooligy system, developed with help from DARPA, could be used with graphics chips and large programmable gate arrays, too.

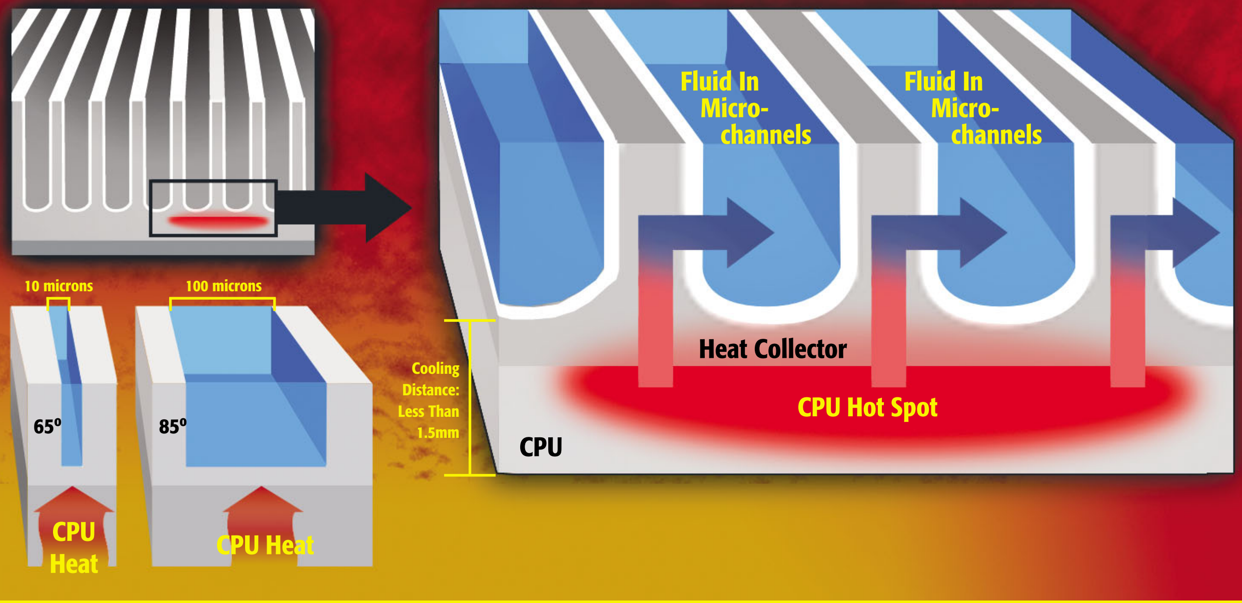

3Com®

The Company That Wired the Network Now Unwires it

Introducing the 3Com® Suite of 802.11 Wireless Solutions

Imagine untethered network and Internet access from anywhere in your home or office. 3Com® brings freedom of choice to wireless with one of the most comprehensive wireless networking suites on the market today.

3Com wireless products come in a variety of price points and are scalable from home and small office to enterprise solutions. Our OfficeConnect® line for home and small offices is among the first to ship with the latest 802.11g standards and new 256-bit Wi-Fi Protected Access (WPA) encryption security.

Our enterprise-class products include an 11a/b/g PC Card with patented XJACK® antenna and dual-mode upgradeable access points that offer you the freedom to mix and match 802.11 a, b, or g standards to meet different coverage and bandwidth needs, and featuring the latest suite of security standards enabling customers to match the level of security to their specific environments.

Best of all, every 3Com wireless product is backed by 30 years of Ethernet networking expertise.

Feel free to unwire your office. Visit www.3com.com/shop today to find the 3Com wireless solution that works best for you.

3Com® OfficeConnect® Wireless 11 a/b/g PC Card, 11g Gateway

3Com® 11a/b/g Wireless PC Card with XJACK® Antenna

3Com® Wireless Access Point 8750
Tri-mode. Upgradeable.

To find a reseller nearest you, go to: www.3com.com/locate_reseller

CDW® Insight® High Tech Low Cost Smart Business PC Connection® PCMall

Copyright © 2003 3Com Corporation. All rights reserved. 3Com, the 3Com logo, OfficeConnect and XJACK are registered trademarks of 3Com Corporation. All other product names may be trademarks or registered trademarks of their respective companies.

E-voting

Soapbox Salvation Or Political Suicide?

Sequoia Voting Systems has several e-voting products, including the Edge touchscreen model.

Think back to Nov. 7, 2000, and the night of the U.S. presidential election (which, for many TV viewers and news junkies, extended into the predawn hours of Nov. 8). Problems with confusing ballots and vote tabulation in the state of Florida—followed by the corresponding legal battle—left the nation waiting for word on an official election winner for more than a month.

Regardless of your personal feelings about the outcome of that race, there's one thing about the 2000 presidential election on which everyone can agree: Voting shouldn't be this hard. In the months after the election, many pundits made bold predictions for future presidential elections in the United States, including the widespread use of new voting techniques, such as Internet voting and electronic polling places. Although full-scale e-voting and Internet voting are still several years away, some voters participating in the 2004 presidential election in a few months will see big changes in the way we vote.

The Paperless Ballot System

E-voting can be a little difficult to define. It's more than computerized voting because computers have been used in voting and elections for dozens of years, especially in the tabulation of votes. Most people would define e-voting as a paperless ballot system, in which vote tabulation takes place automatically. In an e-voting system, voters could use a touchscreen or an ATM-like box at a centralized polling place to record a vote. Other e-voting ideas would have voters record their votes via a secure Internet connection on a computer or on a PDA from almost anywhere.

Voting Systems In Use

Election Data Systems, a political consulting firm, has projected that about 50 million registered voters live in a jurisdiction that will use e-voting systems during the November election.

Voting equipment	Voters*	Percentage
Optical scan	55.7 million	32.2%
E-voting	50 million	28.9%
Punch cards	32.2 million	18.7%
Voting machines	22.2 million	12.8%
Paper	1 million	0.6%
Mixed methods	11.7 million	6.8%
Total	172.8 million	100%

*(*NOTE: Number of voters is equal to those registered for the 2002 general elections.*)

SOURCE: ELECTION DATA SERVICES

Vote-Counting Technologies

Even though some jurisdictions–especially those in small population areas–still use a pencil, a paper ballot, and an X and still count votes by hand, most of today's elections involve use of a computer at some point. E-voting machines give voters direct interaction with a computer, but even in those jurisdictions that don't use e-voting technologies, computers will still play a key role in tabulating votes in November.

Optical scanners. In optical scanning technology, a computer scans paper ballots, searching for marks that a voter makes on the ballot with pen or pencil. The optical scanner then counts the ballots automatically.

Punch cards. Here the voter typically puts a hole in a paper or thin cardboard ballot to indicate the candidate of choice. A computerized counting machine then tallies the ballots by reading the holes.

Voting machines. Although voting machines have been around for more than a century, they've evolved significantly during that time. In most cases, voters push buttons or pull levers to indicate their choices. A computer then calculates each of the votes. Voting machines are being phased out as part of HAVA (Help America Vote Act).

None of these voting technologies is perfect, however. Computerized counting machines can have difficulty correctly counting the ballots when a mark isn't perfect or when a chad doesn't completely detach, blocking the hole. For example, in the March primary election in California, an optical scanning machine misread more than 6,600 absentee ballots because of a mistake by an election worker in setting up the machine.

Some researchers say a counting machine could scan the same 1,000 ballots 10 different times and come up with 10 different results because of inconsistencies in the ballots. Voting machines can be expensive and are subject to breakdowns. Voting machines also typically don't have a backup system, leaving them subject to the same types of fraud and manipulation that critics are pointing out in e-voting machines. Even something as simple as putting an X on a ballot can lead to human error in counting the results.

E-voting isn't perfect, either, but many experts agree that it would prevent or eliminate many of the human errors that voters and election officials currently make. Others say without a paper backup system, e-voting is doomed to constant scrutiny and doubt over the authenticity of results.

But just as there are two candidates in presidential elections, there are two sides to the e-voting story. On the positive side, most analysts expect voter participation would increase significantly if voters had the option to use the Internet in an election. In addition, e-voting should speed the process of counting votes and make the counting process more accurate. Finally, e-voting might prevent the confusion over ballot formats that plagued Florida voters in 2000. For example, the computer could catch a balloting error, such as inadvertently voting for more than one candidate, before accepting the ballot.

Lack of trust. On the negative side, the biggest obstacle for e-voting is a lack of trust among some voters, watchdog groups, and government officials. With widespread concerns about verifying e-voting results among a variety of people, e-voting probably has more hurdles to cross regarding public relations than technological advances.

Hack The E-vote

One of the main concerns with e-voting is how to make it secure from hacks. The diagram below outlines a potential problem and its possible solution.

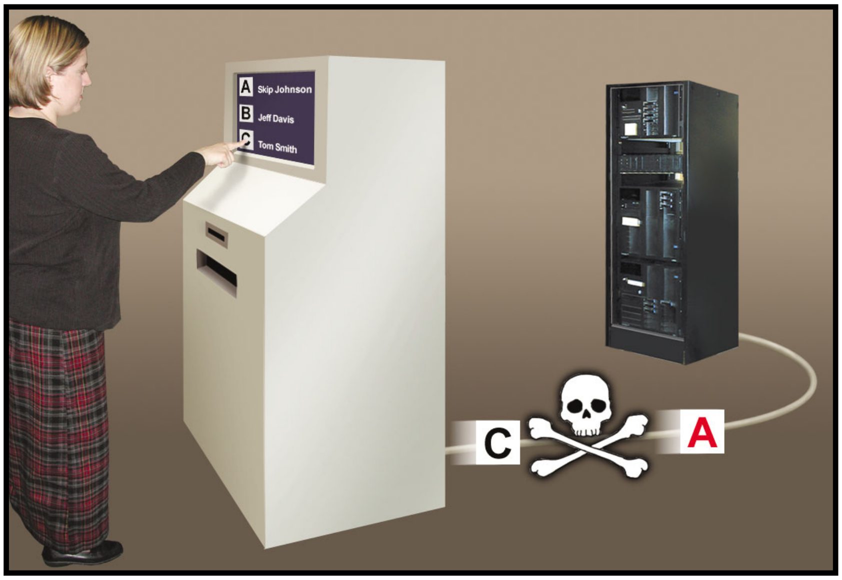

Election officials worry that a hacker could infiltrate the e-voting software, causing the software to record the votes incorrectly. Because the voting system is completely digital, it's possible no one would know about the manipulation. Inadvertent coding errors in the software could cause the same problem, leading to incorrect election results.

To work around this problem, many states are considering including paper receipts as part of the e-voting process. After the voter makes his choices, the e-voting machine would print a ticket showing the voter's choices. The voter would double-check the accuracy of the recorded votes and then present the ticket to an election official, who would securely store it. If a recount is necessary or if election fraud is suspected, the paper tickets would serve as the recount mechanism.

Thank you for voting.

Watchdog groups have been very vocal about their concerns with e-voting. Although some of the concerns listed below are unlikely to be a problem, the simple fact that they could occur is enough to give these groups a strong following. News reports about embarrassing mistakes with e-voting systems and odd results have fueled the fire.

Foremost among the problems is security. Critics of e-voting say there's no way to secure e-voting systems to prevent the possibility of fraud. One hacker or a technically savvy insider could alter the results of an election by figuring out a way to vote hundreds of times or by altering votes. These concerns are causing the biggest PR headaches for those who manufacture the e-voting machines and those who support e-voting.

Other concerns. Some people worry about anonymity. How can Internet voters—who would have to use a PIN or some other type of ID to enter their vote—be sure their votes remain anonymous? Critics also worry that those who aren't familiar with computers could have difficulty voting or may be so intimidated that they stay away from the polls.

Also, with no paper ballots, it may be tough or impossible for a candidate to request a recount in a tight election. Votes would only exist digitally, leaving election officials with no backup option if they later discover an e-voting machine has malfunctioned all day. With machines that count paper ballots, election officials at least have a backup option if voting results are questioned.

Others are questioning the backgrounds of those who create the e-voting machines and software. Using proprietary software could lead to accusations of fraud because members of the public cannot view the software code. Jurisdictions in Australia have addressed this problem by creating an open-source code to run their e-voting machines.

Some voting jurisdictions have one final problem: cost. Any e-voting system is going to cost the government some money for new equipment and software-money that many governments may not have in a time of budget crunches, especially in areas where the voting population is too small to justify the expense. This problem could create a gap in the voting technologies used from precinct to precinct, making the elections unfair.

After the initial expense, however, e-voting should end up saving governments money in terms of the annual costs of running elections. Governments won't have the recurring expenses of printing paper ballots, and ballot tabulation would be far easier and more cost effective.

E-voting Results

Several thousand voters in California and Arizona experimented with e-voting in the 2000 presidential election. E-voting for governmental elections has been popular in Europe and other countries, too. E-voting is already widely used in Belgium and Switzerland, while e-voting trials are occurring in the United Kingdom, France, Spain, Denmark, and Ireland.

After the U.S. government pledged in 2002 to make nearly $4 billion available to enhance e-voting technologies, as many as 34 states developed plans to use e-voting for governmental elections (at least in certain jurisdictions). However, much of that money has yet to trickle down to the states, leading to problems with finding money to train election site workers.

Training trials. Although many jurisdictions reported minor problems or no problems with e-voting technologies during primaries early in 2004, a lack of training for site workers was apparently partly to blame for problems in Georgia and Maryland, as well as some California jurisdictions during the states' March primary elections. In California, about 40% of voters in the primary election used touchscreen machines to record their votes. Problems there included broken machines that prevented some from voting, incorrect ballots loaded onto the machines, uncertified software loaded onto the machines, and slow tabulation of votes. The problems led some California state senators to ask that the machines not be used in the November general elections. However, such requests are unrealistic at this point because of plans already in place.

The Department of Defense had hoped to allow soldiers and other U.S. citizens living overseas to vote electronically in 2004,

E-Voting In Action

Centralized Location

One of the most popular e-voting options from a centralized polling place involves selecting candidates using an ATM-like machine.

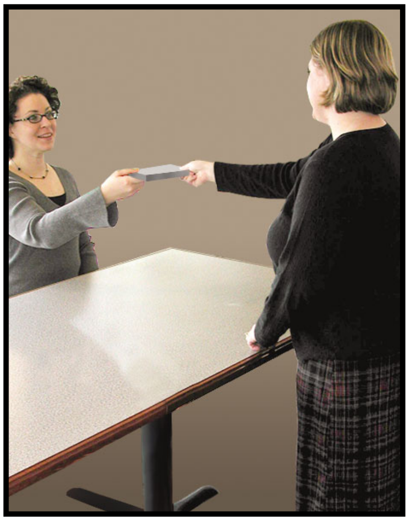

Instead of receiving a paper ballot, the voter would receive a cartridge on which to record her votes electronically. The cartridge would contain voting issues specific to her party affiliation (for voting in a primary election) or for her voting precinct.

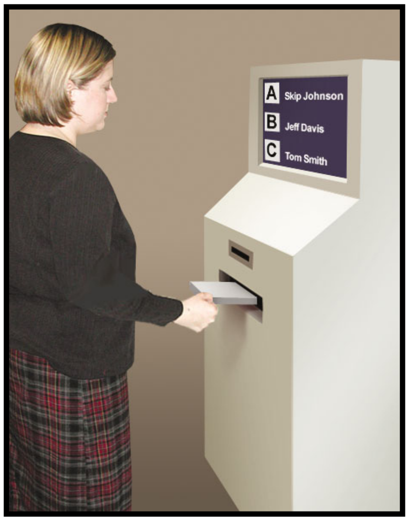

The voter would then enter a voting booth. After entering the cartridge in the machine, she could select candidates from a list that pops up on the screen using buttons along the side of the screen. In some instances, the voter could push an information button, allowing her to see a photo of the candidate or a message from the candidate.

After voting, the voter would remove the cartridge and return it to a worker, who would insert it in a counting machine. The machine would read the information and then send a copy via a secure Internet connection to the election office computer, which would tally the votes in real-time. The results would remain saved on the polling place machine, serving as a backup copy. Finally, the polling place machine would reset the cartridge, readying it for the next voter.

but the DoD scrapped its plans because of potential glitches in the system.

Paper backups. Concerns and problems with e-voting machines have led to calls for a paper backup system. In about half of the states that plan to use some form of e-voting in November, election officials and state representatives have made requests for paper backups. However, many states don't think they can have the paper backups functional in time for the election. Also, paper backups would be expensive. In one Missouri county, election officials estimate they'll spend $25 million on e-voting machines; adding a paper backup system would cost at least $12 million more.

Still, it appears most critics of e-voting would be satisfied with some form of paper backup system. Bills introduced in the past few months in both the U.S. House of Representatives and in the U.S. Senate would address this system, although the changes would be too late for the November 2004 elections.

Not Just For Politics

Although there are several concerns about using e-voting exclusively for governmental elections, e-voting does have one key outlet in the United States: nongovernmental elections. Many businesses and educational institutions have hired e-voting firms to run elections for them, either via the Internet or through a company network.

If you own stock or mutual fund investments, you've probably experienced this type of e-voting. One of the most popular Web sites for this type of voting, www.proxyvote.com, allows investors to cast their ballots online.

The Wait Is On

With all of the concerns over e-voting that have cropped up in the past couple of years, it's doubtful that many U.S. voters in the 2004 presidential election will be casting a strictly electronic vote. Many e-voting jurisdictions are now hoping to at least include a paper backup system. Other jurisdictions will stick with the tried-and-true paper ballots and traditional voting machines. In those areas, changes have been made to ensure smoother elections and less confusing ballot layouts.

E-voting technologies certainly hold a lot of promise for the future, but it's clear that concerns about security and other issues will need to be addressed before we can fully employ the technologies. Those of you itching to vote from your home PC may just have to wait a few years.

by Kyle Schurman

For more on e-voting, subscribers can check out www.cpumag.com /cpujun04/evoting.

No Centralized Location

In the future, voters might be able to record votes wirelessly without ever having to visit a polling place.

Using wireless networking technology, a voter could fill out a ballot on her PDA while reading information on the candidates before the election, thereby making a more informed choice. Once she made her choices, she could indicate on the PDA ballot that she's ready to upload her votes.

When the polls open on the day of the election, the PDA would automatically begin searching for a wireless network WAP through which it could securely upload the ballot to the Internet and to the correct election commission jurisdiction. The ballot would contain enough data to identify the voter for purposes of ensuring she votes only once, but the identifying data would be sent separately from the voting data, protecting the voter's anonymity.

The election commission's central computer would send acknowledgment of receiving the ballot to the voter's PDA. The ballot software then would delete the ballot, ensuring the voter's ballot remains secret. Obviously, the same technique would work on a desktop computer connected to the Internet.

LG FLATRON® LCDs.
Available in LG, XLG and XXLG.

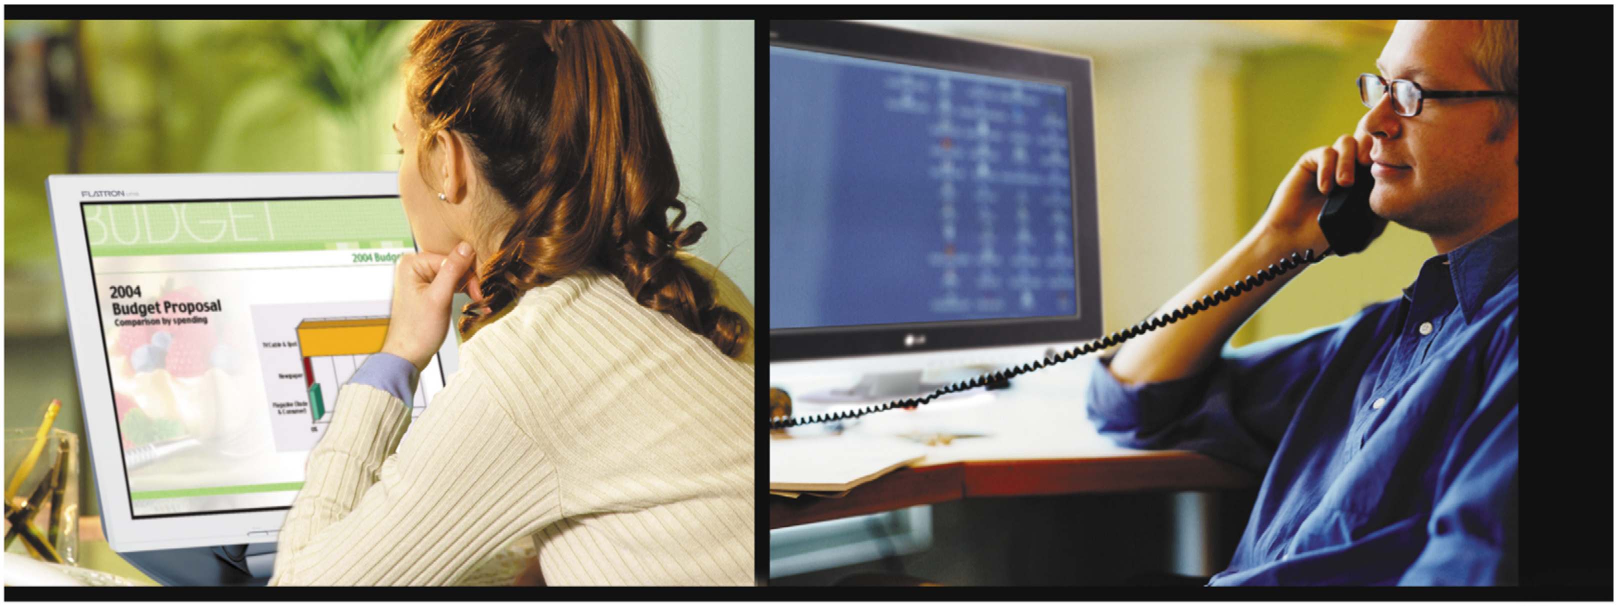

©2004 LG Electronics U.S.A., Inc.

L2320A - 23” Wide-Screen HD/WUXGA LCD Monitor

Available at:

BestBuy.com

THE GLOBAL MARKET
Year II (2002) Proposed Advertising Budget
-By Type & Medium

Because one size does not fit all.

With screen sizes from 15″ to 30″, offering features such as DVI and HD inputs or TV function, there's a perfect LG fit for every home or office. LG LCD monitors offer exceptional quality, innovative features and are all backed by an outstanding 3-year replacement warranty.

The view is better with LG.

www.LGUSA.com

LG
Life's Good

NVIDIA SPEAKS

[The GPU Giant Comes Clean About What Went Wrong, How They Fixed It & Where They're Taking Us Next]

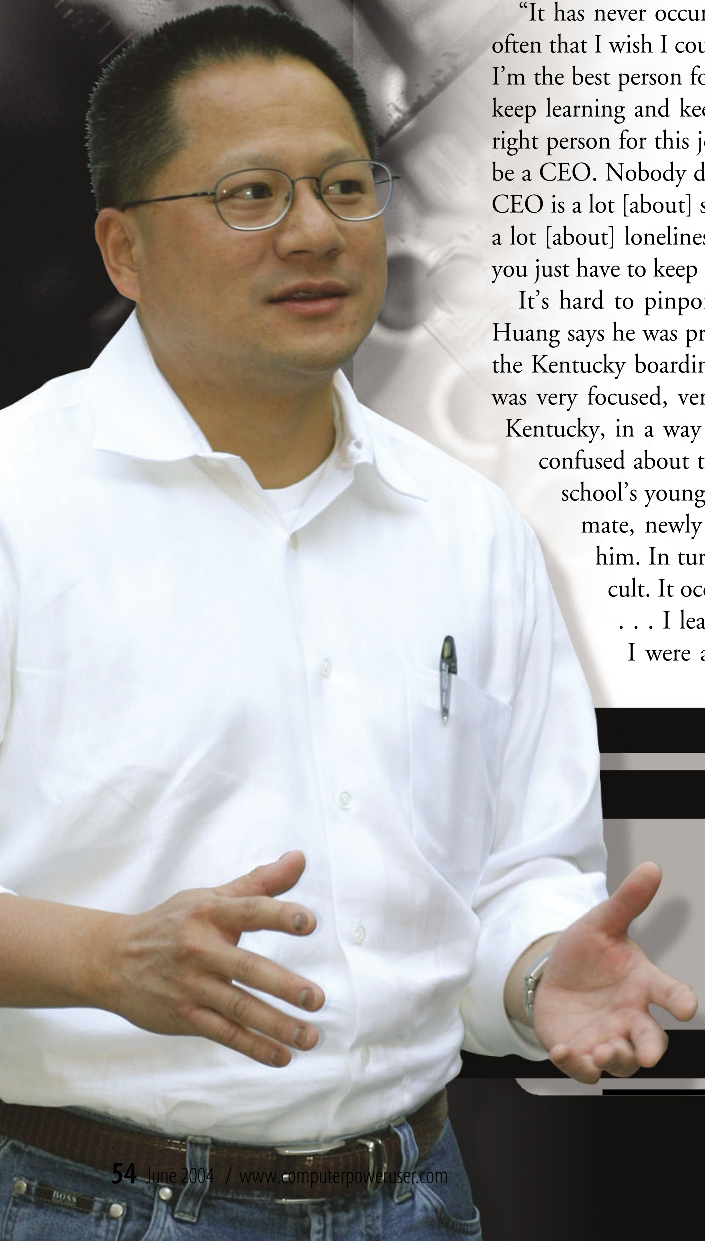

NVIDIA CEO, President, and co-founder Jen-Hsun Huang sips a Diet Mountain Dew (the "drink of CEOs of the 21st century") after spending a good part of the day overseeing company presentations to industry analysts. We are catching him between an interview for *The Wall Street Journal* and a nice, quiet dinner with about 50 international technology journalists flown in for the next day's briefing on NVIDIA's newest graphics chip, NV40. Complicating his image as a tenacious competitor, he comes across as relaxed, occasionally even goofy, with a quiet but infectious enthusiasm. It quickly becomes obvious that he loves his job.

"It has never occurred to me that I shouldn't be the CEO of the company. It occurs to me often that I wish I could do a better job for the company, but I also believe deep in my heart that I'm the best person for the job and I tell myself constantly that I need to keep reaching out and keep learning and keep expanding and keep growing as a CEO or someday I wouldn't be the right person for this job. It is really just a learning process for all CEOs. Nobody was trained to be a CEO. Nobody deserves to be a CEO. Nobody, quite frankly, went to school to be a CEO. CEO is a lot [about] skills but it's a lot [about] leadership, talent, instinct, gut feel, courage; [it's] a lot [about] loneliness. I don't think that there is anything that would train you for it. I think you just have to keep growing as a person."

It's hard to pinpoint when an individual's personality fully develops, but Taiwanese-born Huang says he was pretty far along when he and his brother arrived at Oneida Baptist Institute, the Kentucky boarding school his aunt and uncle had mistaken for a prep school. "As a kid, I was very focused, very determined, very interested in academics and sports. When I went to Kentucky, in a way I was too young to know that it was a reform school, although I wasn't confused about the fact that the kids there were really tough." He was 10 years old. As the school's youngest student, things could have gotten ugly, but Huang's 17-year-old roommate, newly released from custody and sporting a few stab wounds, watched out for him. In turn, Huang taught him to read. "It never occurred to me that that was difficult. It occurred to me that it was exciting and interesting and people treated me well . . . I learned a lot and it taught me a lot of independence. My older brother and I were away from our parents for almost two years." To stay in touch, Huang's

54 NVIDIA CEO: Jen-Hsun Huang
pg 56 Company Core
pg 58 Lessons Learned
pg 61 A Look Ahead
64 NV40: It's Here

parents would send audiotaped messages, which the brothers would record over and send back. Perhaps that is the origin of Huang's interest in technology and its potential to facilitate visceral, emotional experiences. Wherever it came from, he firmly believes that orientation has always set NVIDIA apart from other semiconductor companies.

> “There's not an ounce of me that wants to start another company.”

“When we started our company, the majority of the world's semiconductor companies were founded by semiconductor industry veterans. They built factories and they manufactured chips for the industry. We had the perspective that we were a unique systems and computer company in the sense that the type of products we build deliver an experience; that content, art, and applications are at the core of it and that our technology is really a facilitator of it. We always had the perspective that delivering the experience was our ultimate mission.”

Yet, NVIDIA's road to success was not necessarily a predictable one. “I never thought of myself as a person who was going to start a lot of businesses. I know I'm definitely not a serial entrepreneur.” Huang admits, “There's not an ounce of me that wants to start another company. Not because it's too hard, not because I don't think there are other opportunities, but because that's just not my style. My style is to believe in something, give it everything I have until it doesn't want me any more, for whatever reason. If it weren't for a couple of people I really valued and really believed in, [NVIDIA cofounders] Chris Malachowsky and Curtis [Priem], and the circumstances that surrounded 1993, I don't think I would have started a company. I would have stayed at LSI Logic. I would have worked there until I was running the company. I'm an entrepreneur in the sense that I always have ideas. That's the part of me that is always thinking . . . about ways to do something better, ways to kick it up a notch. I'm always looking for new ideas and I'm never satisfied." This goes beyond grace under pressure and appears to permeate the entire company. "I always believe that the best ideas come out during the worst times, that the most inspired ideas and the most inspired work come during the worst adversity. It is just a fundamental belief of mine. I think best under pressure. I work best under pressure. I am most fired up when we are flat [on] our back, and I always see a silver lining. I really pride myself on being balanced. Things aren't horrible; things aren't perfect. But if you're not focused and if you don't give it everything you have, it's not going to just come your way naturally."

Huang's easy confidence comes off without egotism. "I surround myself with really smart people, making sure that I don't become so hung up on my own ideas that I don't give them the opportunity to show me the right answers, and making sure that [I] keep my guard down so that they're willing to come and tell you when things aren't going well. If you just practice those few skills, and it's counter-CEO, because you're supposed to be the super-smart guy and you're supposed to be the know-it-all guy, but if you can fight that off, you can learn amazing things and people help you, they protect you, they protect the company, and they make you successful.

"There's no secret sauce for having great ideas; the secret sauce for having great ideas is trying many of them. That's the secret sauce."

With only a branching shader for modesty, NVIDIA's Nalu character floats effortlessly while various lighting techniques play with her delicately rendered hair.

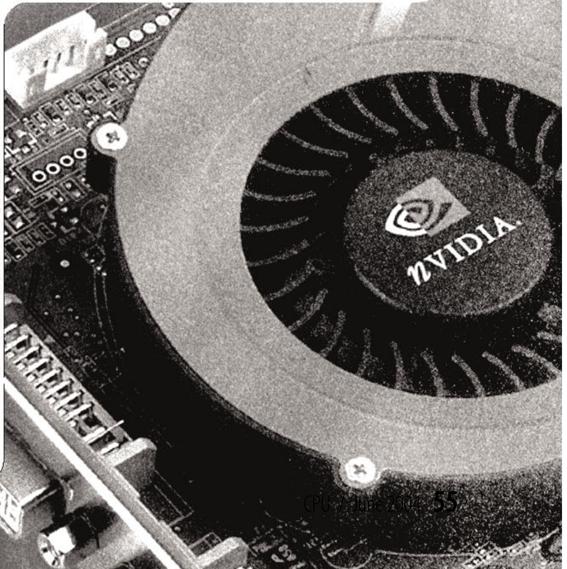

In The Beginning . . .

It speaks volumes about the man that in co-founder and Vice President of Hardware Engineering Chris Malachowsky's official corporate portrait, he is the only NVIDIA executive not wearing a necktie; it didn't occur to him to wear one to the photo shoot, and when they asked if he wanted one, he declined. Extremely well-liked by NVIDIA support staff, he is notorious for his Halloween costumes. (Last year he donned black leather chaps and spiked hair as the ne'r-do-well brother of NVIDIA icons Dawn and Dusk. A snapshot of him in full regalia is still on his official NVIDIA ID pass.) "My objective in life had nothing to do with business. It was just to do interesting stuff and build products people liked." Lest you think he is not a serious person, his desktop motto is this: Success is not so difficult—just bite off more than you can chew and then go out and do it.

According to Malachowsky, the three NVIDIA founders developed a relationship back when he and Curtis Priem (who retired from NVIDIA in 2003) were working at Sun Microsystems, and Jen-Hsun Huang was their silicon vendor at LSI, facilitating production and manufacturing. As the two Sun engineers were moving up in that organization, Huang was transitioning into a business role at LSI, licensing technology and cutting business deals. The original idea for the company came out of a cancelled Sun project. Priem and Malachowsky wanted to design a graphics controller to demonstrate the use of a new DRAM and license it back to the DRAM manufacturer, Samsung.

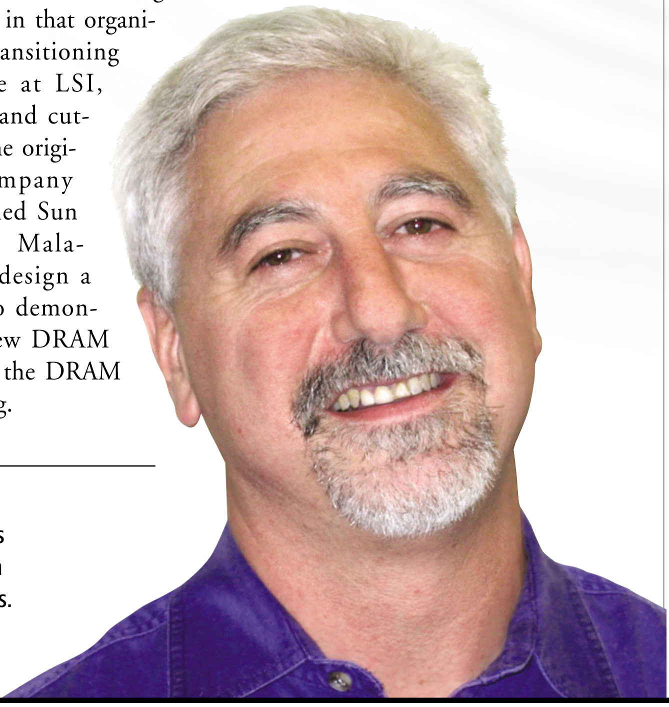

VP and co-founder Chris Malachowsky has a pilot's license and a cabana with multiple plumbing fixtures. Google does not lie.

"Now that sounded like a good plan except that neither of us knew diddly about licensing or dealing with an international company from a business point of view, and we were kind of worried about that. Jen-Hsun was now in this business of licensing technology, and the business end of the relationship grew from there," Malachowsky remembers. "So from day one Jen-Hsun was going to be the business guy. He was the only one of us with any business background, and it was clear he was the CEO of the operation. I was going to be the engineering guy and Curtis was going to do architecture."

Malachowsky sites the mutual respect for each founder's abilities and the clear delineation of their respective roles for the lack of infighting. "It wasn't that we didn't argue a lot about what were the right things to do and the right things to be focusing on, but in the end we would decide. If it was architecture, Curtis would decide. If it was business or strategy, Jen-Hsun would decide. If it was implementation, I would decide and we would go on about our lives. We would trust, even if we didn't agree."

> **["The days of designing chips just based on sweat and effort are gone."]**

Cutting-Edge Tools For Bleeding-Edge Development

What worked early on for NVIDIA is no longer a viable way to make semiconductors. As Malachowsky puts it, "The days of designing chips just based on sweat and effort are gone." One crucial element of groundbreaking innovation now is being able to pioneer the use of new tools and methodologies, many of which are being developed and tested on the fly. "The chips I've worked on have always been the largest in the industry at the time I was working on them. Therefore the tools and approaches that were stable, well tested, and battle proven weren't appropriate. They didn't have the size, capacity, or performance, or they didn't address whatever the new technology was. After we get done plowing through a large design on a new technology, in subsequent years that's the technology that everyone will make use of. But at the time we're doing it, there are only a few companies out there on the bleeding edge." Thus, a mutually supportive partnership with independent tool developers becomes essential and Malachowsky is not shy about speaking out on behalf of these guys. About half of his NVIDIA-related Google search results are for speeches or endorsements he has made for various chip design/verification tool companies. "We want to be supportive of their efforts so that there is a financial gain for them in trying to keep up with us. Ultimately there is nothing terribly proprietary about these tools. Somebody has to do the effort. It makes no sense for us to develop them ourselves, even if we knew how."

But Malachowsky does not credit technology alone with NVIDIA's rise to dominance. "What I didn't appreciate prior to NVIDIA was the value of all the strategic stuff: the business and sales and marketing stuff." He always thought they had good technology. "Our engineering team was second to none. But if you hang your hat on just one thing, and the typical thing is the product, [that] is not sufficient in this market for business success."

This fish-eye lens shot distorts NVIDIA's Santa Clara, Calif., headquarters considerably. The logo really isn't bigger than the moon.

Core Business & Corporate Survival

Huang's passion for graphics only enhances his view of the larger picture, "Our core business started out being graphics, and 3D graphics was very important to me and to the company, but at the core of it, I've always seen NVIDIA in a very simplistic way: I want us to be and I believe we can be one of the most influential technology companies in the world. If you say that when you are zero billion dollars large and three people large, it sounds kind of silly . . . but when you are 2,000 people large, it's starting to sound possible." This kind of talk may come across as corny (or worse) on the printed page, but having spoken with long-term NVIDIA employees who used to roll their eyes at this type of company cheerleading but now are seeing the vision too, it can be pretty compelling in person. More than once, the image of Bill Gates, spinning his vision for the PC industry back in the early 1980s, came to mind. "Being able to add value to the industry and being able to change people's lives has always been at the core of what we believed. Now, is it central around graphics? In a way, I hope so, because it is such a visceral, such an emotional technology, and actually makes a connection with you. So from that perspective, it's exciting. But I don't need it to be, because I think that is one of the reasons we called our company 'NVIDIA' vs. '3-D-F-X.' I think NVIDIA can do a lot more than just graphics and we do a lot more than graphics today."

Malachowsky's observations on corporate survival are even more to the point: "No one has been fired for having a product be late [or] misfunction. We ask everybody to work as hard as you can, use your best judgment, decide what you think is the right thing, and run with it. Make sure everybody knows where we're going and if it doesn't work out, be the first one to say, 'I need help.' One of the things that we've enabled here is that every group is chartered with, is expected to, and is waiting to carry the day. We are considered a hardware company, but if the chip is late, well shoot, software better write a better driver, it better look fresh. If software can't eke out more performance, then maybe operations needs to get a chip build that's faster. If they can't do it, marketing better make up some feature and tout the crud out of it such that it seems fresh. If that doesn't work maybe sales has to come up with a new pricing strategy. . . . The tactics, the strategies, [of] how you relate to your customers, how you deal when the schedule is late, what do you do when the product costs too much, when it doesn't yield right, when competitive threats come from here and there, when the press is friendly or not friendly? Throwing up your hands is not the right answer, and one of the strengths we've always had is, the larger the crisis, the crisper the thinking. In general, most things can be turned around either with effort or insight or clever thinking. The way we've responded to calamities and challenges actually defined us and helped ensure that we survived another day."

This may never have been more evident than in the past 18 months, when the very late arrival and somewhat disappointing performance of NVIDIA's NV3x family of graphics chips kept engineers fixing, marketing folks spinning, and customers waiting. Huang's 2004 Editor's Day presentation focused primarily on the many lessons learned from NV3x, what has been carried forward into NV4x, and how NVIDIA has been transformed by the fallout.

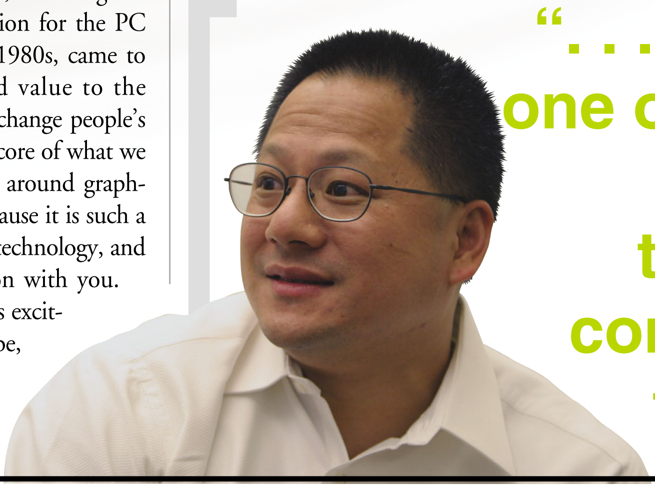

> ". . . we can be one of the most influential technology companies in the world."

> "We took a lot of flack for NV30. We made some pretty big mistakes, but we also did some really great things."

What Went Wrong

With NV3x behind him, Huang is more candid than he could afford to be during the last few product cycles. When NV30 didn't appear for Christmas 2002, despite assurances from Huang and others, it cost NVIDIA a lot more than revenue; it cost credibility. When asked why analysts should believe his claims regarding the current NV40 chip, Huang simply replied that the current product is not eight months late.

The short list of NV3x technical problems isn't actually very short. Some are manufacturing-related, such as opting for a largely untested (for graphics) 0.13-micron process, which certainly played a large role in delayed time to market. Others stem from inappropriate technology choices, such as the big math machine VLIW architecture, which was a nifty idea but so unworkable in practice it was scrapped altogether for NV4x in favor of "wider is better" parallelism.

Many items, such as FP32 and the run-time compiler development, appear to have been mistimed. FP32 is definitely the way to go eventually, and even though Huang asserts strongly that it was the right choice for NV3x, it is a hard sell in light of how expensive it was (Huang says it cost NVIDIA a quarter of a billion dollars in operating profit last year) and how little developers could make use of it within a reasonable time frame. Had early communication with Microsoft been completely clear regarding DirectX 9 floating-point precision requirements, NVIDIA may have chosen the lower 24-bit precision ATI continues to employ. The run-time compiler, quite useful for a programmable part, was not developed simultaneously with the architecture, perhaps was not staffed with seasoned compiler programmers early enough, and was ultimately diverted into fixing specific programmability shortcomings of the architecture rather than focused on overall performance and ease of use.

On top of the technology challenges, there seems to have been a fundamental flaw in NVIDIA's project management method. Though not directly responsible for NV3x or 4x, as a veteran of earlier NVIDIA chip engineering cycles, Malachowsky explained that from the company's inception, there was never much in the way of hierarchy. "Whoever was responsible for the chip handled all aspects of the design and the verification and the implementation. He coordinated with the software guys, the board guys, and the systems guys. . . . There was one guy, and everything was under him." Over time, as projects became more complex, this approach apparently became difficult, even unworkable. As he observed in a presentation in June 2002, "Project scope has just about outstretched a single person's ability to understand it all and be enough of an expert to lead it appropriately." By then it would have been too late to completely revamp the management of project NV30.

Created to demonstrate high dynamic range lighting effects using 16-bit floating-point math, Timbury is NVIDIA's "first fantasy character." Apparently, pixies and mermaids are "real."

Lessons Learned

The rightness of specific technology choices will fluctuate from project to project, and Huang has a long list of the good and bad outcomes, benefits, pitfalls, and hard-won lessons from the NV3x mistakes. When faced with lemons, he definitely comes across as a "make lemonade" kind of leader: unafraid to scrutinize errors, determined not to repeat them, and always looking for that silver lining. He says, "Two or three years ago, NVIDIA had the vision that programmable shading was going to change real-time graphics." Until that time, Huang describes, the technology of 3D graphics had been mostly about the evolution of military and flight simulators with everything based on geometry and textures, which, even if painted by a wonderfully talented artist, are created at a time before the game is being played. "Texture mapping, fundamentally, is a sophisticated decal, and there is only so much you can do with a decal. Programmable shading is the right direction because video games are at the intersection of science, technology, and art. It's a complicated thing. It's hard for an engineer to think like an artist, for an artist to think like a physicist. Yet in order to build a great video game, you need all of those skills. So what we've really done with programmable shading and programmability is put the artistic pallet in the hands of the content developers instead of capturing it on our chip. We're engineers. We can draw the polygons and [apply] texture maps, but beyond that, how would we know what [is] the next beautiful effect or the next artistic look?"

Huang's goal of having NV30 take 3D graphics to the next level fell somewhat short. "NV30 was designed to be programmable, but it was actually not a very programmable chip. It was great at legacy applications, which was kind of a weird thing. [Developers] just loved DX9, they went to it right away and that was good and [NVIDIA's high-level shading language] Cg pushed that along, but talk about irony. We told everybody, 'This is the future and we ought to move to programmability.' And when the time came and we built our own chip, we went 'Um! This didn't turn out the way we thought.' "

Yet Huang sites the unbalanced and nonorthogonal nature of the NV3x architecture and the difficulty developers had programming for it as the impetus for NVIDIA to write the real-time optimizing compiler. "Compilability will always imply an elegant architecture to program to but you need to invest deeply in compiler technology. Give me an example of a great microprocessor company who doesn't believe investing in compiler technology is necessary."

His emphasis on including forward-looking features is also quite emphatic, "If we want to continue to push content development to the next level, we need to make sure that we implement features that run at full limits, full rates, even though no games access them today. I have no idea how many times we were asked the question, 'Sure, you've got these great features, but how many games are available today?' I now have found the perfect answer for that. I'm perfectly comfortable saying it, and I said it yesterday to a bunch of financial analysts: None! So sue me! If we don't implement the features, the games will never use them. It's back to the chicken or the egg issue again, and that's OK. If you want to be an innovator, if you want to push the market forward, you're going to have to put features in your chip that nobody is using today. If you're working closely with developers and the industry, those features will be used."

The Dustbuster. "Nobody likes to hear harsh criticism. But we have to hear it." With high chip temps a very real problem, NV30's variable speed cooling solution was crowned "The Dustbuster" by some reviewers. Huang jokes about his personal contribution. "I made a mistake during the launch of the first GeForce FX. I had really pushed us for a very powerful fan. I didn't actually ask [for] it to be loud; I asked for it to be very powerful. It turns out that powerful typically equals loud on fans. So here was my idea: It's called 'Silent Running,' ironically, and depending on what applications you are running and depending on the workload, we would control the fan electronically and adjust the RPM, and it would dynamically apply the amount of thermal management on the chip depending on the workload. Well, it turned out that the fan was really loud. So when 3D turned on when you were playing a game, it got super loud, it was sufficiently annoying. Of course, everyone knew it was my idea, and not a single person would let me off the hook on it. They know that I don't mind when they tease me so they made a whole marketing video. It's hysterical." (Download the video at users.pandora.be/darkt/NVIDIAFlowFX_video.zip.)

The lemonade part quickly follows, "Well, look where that technology went. It turns out that on Apple G5, all the fans are electronically adjusted so that when they use demanding applications, the fan turns on more, and when they're just browsing the Web, it [goes] silent. More and more people are using this technology. We're using it now all over the place, in laptops, in some of our other desktop products. The point is that even good ideas can be taken to an extreme, but if you learn something from it and you adjust those ideas and refine them, [they] can become great ideas. So we don't mind laughing at ourselves, and it hasn't discouraged me from taking calculated risks and offering more ideas to the marketing guys, but sometimes one of them will say, 'This sounds like one of those fan ideas from Jen-Hsun.'"

Huang summarizes, "So, we had a speed bump. But NV40, NV50, and NV60 will show very clearly we pointed the industry in the right direction. We didn't capitalize on it ourselves right away, but that's OK. We're going to capitalize on it now."

> **" 'Sure, you've got these great features, but how many games are available today?' I now have found the perfect answer for that . . . None! So sue me!"**

Moving Forward

Overcoming inertia is a challenge, change is hard work, and success breeds complacency, even for engineers. Malachowsky may have felt earlier than most that all was not well in NVIDIA-land. "When you are on top of the world as we'd been prior to the 3x generation, it's hard to make fundamental changes and improvements. You can tell someone to do better and they say, 'Sure. Stocks are going up here, products are selling out over here, and everybody loves us. I'm not really that bad off. I'll deal with improving later.' But as soon as you feel challenged, then you can say, 'Let's be introspective.' We're not perfect. Even the things we thought we were world-class at, there's another level. Let's go find it."

In order to survive, NVIDIA as a whole had to evolve past old habits and methods. "As in almost all of our experiences in the past, what could have killed us made us better." Huang admits, "NV30 changed the way we organized our company. That's how seriously we look at mistakes. We are willing to make fundamental changes at our core. So architecture is now done differently; we work with developers differently; we work with Microsoft differently; we engineer our chips differently; it's just really, really profound what we've done, and we did all that in a year."

"The 3x family challenges were exactly the impetus [needed] for change," explains Malachowsky, "Nothing about NV4x execution is an accident. The plans, the approaches, were all rather deliberate and aimed at improving anywhere we had a weakness. We've gone toward a much more traditional approach to project management. . . . We've now added some more structure. There is somebody responsible for verification, somebody responsible for implementation, somebody responsible for the different facets, and they really are responsible in that they have delivery schedules. [We have] team leads that are the technical drivers, but now we have more management hierarchy. Coming at it from where we were, we're still much lighter than the traditional engineer organization [but now] the guy at the top is not dealing with every issue simultaneously. . . . So the 4x family and its deployment, execution, architecture, planning for bring up, planning for first production, the software driver, the performance, all of that was deliberate. Not everything worked out exactly the way even those plans went, but we were proactive and thoughtful about just about every aspect having to do with taking that product from conception to the marketplace. Shipping first metal on a 200-something million transistor chip is pretty good, and the performance we're getting out of the box is really great."

Huang agrees, "We now have created an architecture that is much more programmable so that NV40 just runs beautifully out of the box. . . . We've made a machine that is so orthogonal, it just looks beautiful on paper, it's easy to program, developers know exactly how to program it and can almost predict what the behavior will be. . . . One of the good things that we learned on NV3X [is that] the architecture that shows up for the enthusiast or the workstation marketplace and the architecture that we offer you for the entry-level product that you can go down to 7-11 and get out of a Coke machine, that architecture ought to be identically the same. The only difference is performance. The features, the architecture, the programming model ought to be the same. Scalability is very important." This is a large leap from the GeForce4 MX days when previous generation products were branded as current generation technology, confusing customers and frustrating retailers (while giving *CPU* columnists a lot of carping material).

An especially exciting area of increased activity is the involvement of game developers earlier in the design process. "The first hundred NV40s that left our company did not go to OEM customers, did not go to our add-in card partners. The first hundred NV40s that left our company, not even an exception of one, not one came to my house, went to developers. The reason for that is very clear: This is about content." Huang said in his editor day presentation, "For the first time in our history, NV40 was developed at the architectural level with developers. We brought them into the architectural definition phase. We also learned that we need to provide tools to the developers so they can figure out how to program best to our architecture." NVIDIA has ramped up developer support with new tools that allow programmers and artists to sit in front of the same screen and analyze the performance hit that proposed art or effects will have on different NVIDIA platforms in real-time.

> **"Nothing about NV4x execution is an accident."**

Featuring an adaptively tessellated ocean topped with dynamically computed foam, "explosion, smoke, and fire are present as the ship is destroyed" in the oddly named "Clear Sailing" demo.

IT'S SHOWTIME

Talk about box office hits. Verbatim's new DigitalMovie™ DVD recordable media is so hot you can almost smell the popcorn. Available as DVD-R or DVD+R, this 4X speed DVD media looks – even feels – like a Hollywood movie reel. But don't think you need a director's chair to use them. These mini-reels are perfect for home movies, business presentations, transferring files – anytime you want to add a little fun to the show. And of course, DigitalMovie DVD has the same unsurpassed performance and reliability you've come to expect from Verbatim. So make your next show a feature presentation. DigitalMovie DVD from Verbatim. It's the reel thing. For more information, contact us at 800-421-4188 or online at www.verbatim.com.

Verbatim.

PUT YOUR WORLD ON VERBATIM

© Verbatim Corporation 2004

Beyond NV40

There is no disputing how hard NVIDIA worked to recover from the spectacularly impressive triple Salchow face plant of NV30. The follow-on NV35 and NV38 chips proved that the company could refocus under pressure, and by all indications, NV40 is everything its creators hoped it would be. (See benchmark analysis and technology details of the GeForce 6800 Ultra in our hands-on preview on page 64.) In a diligent effort to reinvent itself, NVIDIA has not balked internally at admitting mistakes, rethinking fundamental technologies, and reorganizing production methods, all of it, reportedly, without the customary pointing of fingers and rolling of heads. If only NASA could do that.

But is it enough? Generally speaking, the market does not reward effort, it rewards results. In the past few years, NVIDIA has held onto most of its large market share, despite not having the best high-end desktop technology. Brand loyalty and solid marketing, along with a product that could compete (with a bit of fudging here or there) in the same arena as the top performer, kept things going along fairly well. Should the current tussle between NV40 and ATI's R420 go well or badly for NVIDIA, it could have a big impact on sales of its desktop graphics products, but the future of NVIDIA is based on more than just making pretty pixels for your favorite run/kill/die game.

A New Era

NVIDIA's wider market focus starts with what Huang calls the era of "Networked Digital Media." The idea is that the proliferation of digital content spreading across the PC, consumer electronics, and handheld devices through shared standards and formats, together with the arrival of the necessary infrastructure (ubiquitous broadband, cheap storage, and really big LCD TVs) to distribute, archive, and display high-quality media is creating a massive opportunity. According to Huang, NVIDIA has something to offer the Networked Digital Media Era in three key technology areas. The Network, Media, and Connectivity Processing arena is addressed by NVIDIA's nForce MCP (Media and Communications Processor) technology, which is set to capitalize on the transition to 64-bit computing. The Digital Media Software category covers all of the separate technologies contained in NVIDIA's ForceWare drivers, including a high-quality multimedia server, game profiles, multimonitor and HDTV support, system utilities, RAID disc management, an audio mixer, and firewall network security.

Graphics and Video Processing is the largest category, encompassing all of NVIDIA's GPUs and its advanced capabilities in programmable shading, high-definition video processing, and image processing. Although the competition has been incorporating video processing for several generations, even calling their product a VPU (Visual Processing Unit), Huang's view for NVIDIA technology is a very long one: "NV4x is our first-generation architecture where we really reached out and thought of programmability and thought of the GPU as a digital media processor instead of just as a graphics processor. It is a transformational idea for our company, and I believe that in 10 years, our GPUs will be used to drive all kinds of stuff, from car navigation systems to projectors to digital televisions to you name it."

Branding & Licensing

A unified marketing message is essential, and there are currently six brands in the NVIDIA store, each with a distinct purpose and market: GeForce-branded GPUs on the desktop, GeForce Go for mobile, Quadro in workstations, GoForce processors for the handheld market, the nForce brand MCP southbridge for motherboards, and the ForceWare branded unified software environment to pull it all together.

NVIDIA is likely to continue its habit of entering into cross licensing agreements with other computer graphics technology companies, which Huang

NVIDIA Brands

NVIDIA ForceWare™

NVIDIA GoForce™

NVIDIA nForce™3

NVIDIA GeForce™

NVIDIA GeForce™ Go

NVIDIA Quadro®

sees as a liberating force allowing engineers to take best design paths instead of cumbersome workarounds. “What a cross license does is enable you the freedom to innovate. If IBM would like to use our technologies so that they can build whatever computers they want, terrific. If Silicon Graphics would like the same thing, terrific. Evans & Sutherland? Terrific. I happen to believe that intellectual property is used defensively and used wisely to enable you the luxury to innovate in the marketplace. The thing about innovation is that all good ideas are going to be copied. So unless you want to take on everybody in the world and spend the rest of your life in court, your best defense is to keep innovating.”

Handheld Devices

NVIDIA is ramping up graphics processors for handheld devices very quickly. Last year’s acquisition of MediaQ brought in some advanced ultra-low-power processor technology and the engineers that created it. Now integrated with NVIDIA’s small in-house team under a new vice president and general manager of Handheld, the second-generation handheld product, with an incorporated 3.2-megapixel camera, will be released later this year. NVIDIA is also allocating resources for handheld content developer support. (See “Forward Slash” on page 89 for details.) Huang believes NVIDIA’s long experience processing pixels for a demanding audience will enable it to compete favorably in this new area. “We will be able to add value [because] computing the pixels for cell phones is arguably even harder [than for PC]. You have such a small footprint, whether it’s memory or bandwidth or display size, that you have to take extraordinary care . . . the video that’s coming to you is in such a densely packed and compressed [format] that you have to do a lot of media processing to it to finally deliver the beautiful pixel. In the area of 3D, because you don’t have that many pixels to start with and because it’s so close to your face, you better make sure that every single pixel is highly cared for and comes across beautiful and delicate, otherwise it will look grainy and ugly.”

> “Never underestimate the ingenuity and cleverness of someone who has a goal in life and finds himself slightly challenged.”

Growth potential in this area is huge, and Huang says it is time to “deliver the type of experiences that we imagine a great cell phone would have, whether it’s making a video phone call, watching CNN, using it to listen to music, or playing a small game while you’re sitting on a plane. We believe that the time is finally here.” Again, there is also the long vision, “Over the course of the next five to 10 years, flexible displays are really going to change the way we think about handhelds. These mechanical cell phone-like devices with buttons, that’s going to be the past. Your cell phone is just going to be a single semitranslucent sheet of what appears to look like plastic. It’s going to be bendable. You’re going to touch it, and it will turn on. It will render the buttons instead of having real buttons. It will be very, very high resolution, and it will be beautiful.”

Passionate Pragmatism

In speaking with NVIDIA executives and employees for this story, a common description of the company and its CEO was “pragmatic.” But at the crux of NVIDIA’s corporate culture is a fundamental contradiction: Bleeding-edge innovation is driven not by pragmatism but by passion. It is much more practical to let someone else make the enormous investment in various R&D paths (some that create new markets and some that fail miserably) and simply stand by and observe which ones actually work and how they are utilized before pouncing on the market opportunities. It is far more sensible to be just behind the curve of change if you don’t want to end up under its wheels. Pointing any industry in the right direction is risky. Leading the way with technology before there is a market for it is risky. Passion is risky.

But passion may just be the best game in town. As Chris Malachowsky puts it, “Never underestimate the ingenuity and cleverness of someone who has a goal in life and finds himself slightly challenged. Have you seen the Unreal engine? That thing is gorgeous. There’s no geometry behind it, and that’s pretty impressive, and that’s cleverness. Carmack is clever. Tim Sweeny is clever. These guys and the folks that are like them in the ranks are just smart as whips. They just find a way. The state of the art is continually being pushed.” Though recently elevated to the role of NVIDIA Fellow and no longer on the hook for day-to-day engineering issues, Malachowsky sums it up: “We’re still in the realm of practical. We need to do a practical implementation and developers need to do a practical deployment of their technology. No one’s trying to be ideal. Ideal would probably put us all out of business. But over time we are going to approach ideal.” Words from the wise.

---

by Joan Wood

WEIGH-IN AT THE SCALES

NVIDIA's NV40 Is Here & The Battle Is On

It's hard to forget the follies of NV3x, but with NVIDIA's follow-up NV40 chip, the company is certainly doing its best to successfully leave them behind. The NV40—based on IBM's 0.13-micron process and in a BGA (Ball Grid Array) flip-chip package—is packed with 222 million transistors, almost double the amount yet seen.

In top form, NVIDIA has dropped the "FX" moniker and gone back to an "Ultra" branding for the NV40-based GeForce 6800 Ultra board we've been testing. The engine clock is down from NV38's to 400MHz, but with a monstrous 16 pixel pipelines (up from four previously), 16 texturing units (up from eight), and six vertex shading units (up from three), clock speed is becoming less important. Now that the GDDR3 memory initiative has been paved, the 6800 Ultra utilizes 256MB of Samsung's 550MHz (effectively 1.1GHz) DRAM on a 256-bit memory bus interface, dishing out 35.2GBps worth of bandwidth.

Tomb Raider: Angel Of Darkness

Tomb Raider: AOD is a DX9 title that can make use of Pixel Shader 1.1/1.4/2.0, plus Vertex Shaders 1.1/2.0. We tested the PS2.0 code path, but just because this is a PS2.0 game doesn't mean it's all that pretty. The 6800 Ultra takes the lead in all cases. Somehow at 1,600 x 1,200 x 32 with 4XAA and 8XAF, it dips below the 60fps mark, but when you consider it's still 2.5X faster than the other two boards, you might just question the developer as to where all those polys and shaders went.

Tomb Raider: Angel of Darkness	Radeon 9800XT	GeForce FX 5950	GeForce 6800
1,024 x 768	96.2	59	119.4
1,600 x 1,200	96.2	25.8	119.9
1,024 x 768 4XAA/8XAF	70.3	45.2	103.5
1,280 x 1,024 4XAA/8XAF	44.6	27.7	71.3
1,600 x 1,200 4XAA/8XAF	22.6	19.8	59.2

In terms of theoretical performance, the 6800 Ultra stomps all over the previous generation GPUs with 6.4Gpps (gigapixels per second) pixel fill rate and 6.4Gtps (gigatexels per second) texture fill rate. NVIDIA CEO Jen-Hsun Huang has said, "Wider is better," which in essence is why ATI was better on the previous generation of hardware. The NV3x architecture used a more complex VLIW (very long instruction word) system and was dependent on highly efficient and intelligently written compilers, which proved late in coming. NV4x architecture falls back with more traditional ALUs (Arithmetic Logic Units) using a SIMD/MIMD design.

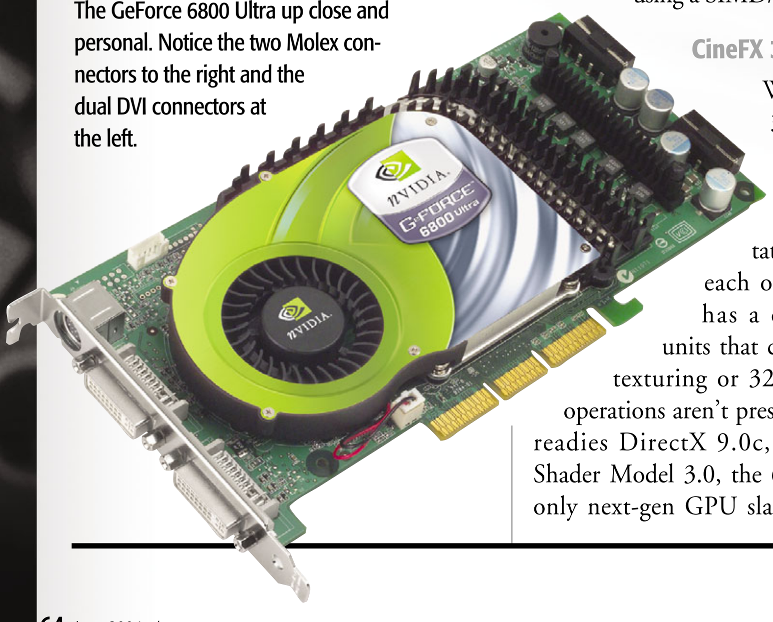

The GeForce 6800 Ultra up close and personal. Notice the two Molex connectors to the right and the dual DVI connectors at the left.

CineFX 3.0

With its CineFX 3.0 architecture, NVIDIA is sticking to its implementation of FP32, and each of the pixel pipes has a couple of these units that can perform 16x1 texturing or 32x0 when texture operations aren't present. As Microsoft readies DirectX 9.0c, which includes Shader Model 3.0, the 6800 Ultra is the only next-gen GPU slated to support it fully. Pixel Shader 3.0 requires the ability for GPUs to take advantage of subroutines, loops and branches, dynamic flow control, and up to 65,535 pixel shader instructions per shader and "full precision" FP32. NV40's pixel shading has additional support for MRTs (multiple render targets), full speed nonpower of two textures with mip-mapping, and centroid sampling AA, and it still allows for FP16 partial precision.

The PS3.0 model also lets developers use position registers, indexed input registers, arbitrary swizzling, gradient instructions, loop count registers, two-sided lighting, unlimited dependent textures, and texture instructions. Temporary registers have been increased to 32 and constant registers to 224. With MRT support it's now possible for the pixel shader to save pixel data into buffers (ATI already does this with its F buffer), which is made available for complex lighting effects post geometry set up and, therefore, doesn't require another pass. Nifty. Each NV40 pipe is outfitted with a second shader unit, so in theory it doubles the traditional capabilities by being able to perform four instructions per pixel per cycle and eight operations per pixel per cycle.

The Vertex Shader model 3.0 extends the instruction limit beyond 2.0's 256 up to 65,535 and introduces vertex texturing with look ups (texture access from within the vertex shader), vertex texture fetching, vertex stream frequency/geometry instancing, and static and dynamic flow control functions, including branching call and return subroutines. VS3.0 allows for vertex processing with textures (aka displacement mapping) and instruction predication, and temporary registers have been increased to 32.

The texture engine can now perform up to 16 texture operations per single pass with FP16 or FP32; there's no FP24 here. Basically this translates into more horsepower and less passes for quicker performance. Additional capabilities include support for nonpower of two textures and sRGB texture format for gamma textures. A close approximation of displacement mapping is now rolled into NV40, which should make conventional bump-mapping fans happy. We have seen some of these capabilities put to good use in the next-gen Unreal engine, and we like it.

Will Shader Model 3.0 take off anytime soon? Probably not, but numerous games, including Far Cry and Painkiller, may get patches to utilize it. How much an upgrade will affect the visual experience depends on if the move is from Shader Model 1.1, which shows very impressive visual differences, or from SM2.0, which may show little, if any, visual upgrades. According to Far Cry developers, who have been experimenting with SM3.0 for a while now, the Crytek engine upgrade from SM1.1 to SM2.0/3.0 that NVIDA used as a demo at the GeForce 6800 launch event is visually identical when run on 2.0 hardware (including ATI's latest) and the NV40 3.0 hardware; however, there may be some differences in performance for things like geometry instancing.

Splinter Cell

This third-person action game was ported over from the Xbox to PC. Although this is a DX8-class title, pixel shading horsepower does seem to be important, and even its PS1.1 implementation stresses previous generation 3D cards, especially in scenes with night vision or water. Multisampling antialiasing techniques don't currently work in Splinter Cell, so we ran the benchmark without FSAA enabled. The GeForce 6800 wins by a significant margin, barely slowing down from 1,280 x 1,024 to 1,600 x 1,200.

Splinter Cell	Radeon 9800XT	GeForce FX 5950	GeForce 6800
1,024 x 768	57.7	41.7	80.6
1,600 x 1,200	41.5	32.1	71.8

Unreal Tournament 2003

UT2003 is a couple of years old now and was designed in the late days of DX7 and early DX8, before programmable Pixel Shaders, which is why the three cards perform remarkably similarly until memory bandwidth and fill rate hit a wall when cranking up the resolutions and turning on AA and AF. UT203's engine has been licensed for many titles and, thus, is still a good test to use. Turning on 4XAA an 8XAF is far from free, even on the 6800, as its frame rate drops from 188.9fps to 100fps at 1,600 x 1,200 x 32.

Unreal Tournament 2003	Radeon 9800XT	GeForce FX 5950	GeForce 6800
1,024 x 768	203.9	206.8	210.9
1,600 x 1,200	102.7	113.6	188.9
1,024 x 768 4XAA/8XAF	127	121.2	181.7
1,280 x 1,024 4XAA/8XAF	84.5	82	134
1,600 x 1,200 4XAA/8XAF	58.5	58.6	100

Unreal Tournament 2004

Being a slight update of the UT2003 engine, UT2004 still doesn't make use of pixel shaders but is slightly more demanding. 3D graphics cards with more memory bandwidth and higher clock speeds do well with this title. There isn't much distance between the boards until you crank up the visual quality. NV40's fill-rate capabilities with its 16 pixel pipes then prove that more horsepower counts. Noting that the frame rate only drops from 152.1fps in non-FSAA AF at 1,024 x 768 x 32 to 125.6fps at 1, 600 x 1,200 x 32 4X FSAA and 8XAF, it doesn't seem like the NV40 is being taxed much. This is a rare case we can remember where the old marketing phrase "anisotropic filtering and antialiasing modes are free on our new hardware" actually seems close to true.

Unreal Tournament 2004	Radeon 9800XT	GeForce FX 5950	GeForce 6800
1,024 x 768	143.6	143.4	152.1
1,600 x 1,200	122	122.5	151
1,024 x 768 4XAA/8XAF	126.2	134.6	150.7
1,280 x 1,024 4XAA/8XAF	91.5	72.7	144.6
1,600 x 1,200 4XAA/8XAF	67.8	67.8	125.6

HPDR

16-bit floating-point textures through the pipeline equal HPDR (High-Precision Dynamic-Range) technology. OpenEXR, ILM's (the "Star Wars" freaks) high-dynamic range format, uses 16-bit floating-point precision written out to the frame buffer instead of just to a graphics processor's registers. HPDR is done in hardware and provides 16-bit floating point formats for blending, shading, texturing, storage, and filtering during the light-transport stage. The tone mapping and color gamma correction stages can have sRGB format applied to them. This is a faster alternative to pixel shader emulation.

Floating-point filtering can improve the quality of images in motion, drastically reducing color banding. We've seen some evidence of this in Half-Life 2 demos already. Floating-point texture capabilities and frame buffer blending give detail to such special effects as motion blur and explosions, and a new rotated-grid technique for antialiasing helps attack those jaggies.

Intellisample 3.0 & Rotated Grid Anti-aliasing

Antialiasing is now a feature that anyone buying a high-end 3D card expects to be able to use. NVIDIA has gentrified its AA engine with the addition of rotational grid AA, which has been a feature on ATI's previous generation chips. It's certainly an improvement on the NV3x, which went with a four-subpixel routine sampled in a 2- x 2-square grid. NV40 uses a diamond-shaped grid for each pixel, using samples closer to the original's pixel plots, giving further color precision at the edge of the polygon, smoothing out the jaggies. Subjectively, the visual quality is much improved. With the new power, plus improved algorithms, 4XAA (antialiasing) and 8XAF (anisotropic filtering) now work properly without dropping down to "brilinear" (except during UT2004) and also have improved performance.

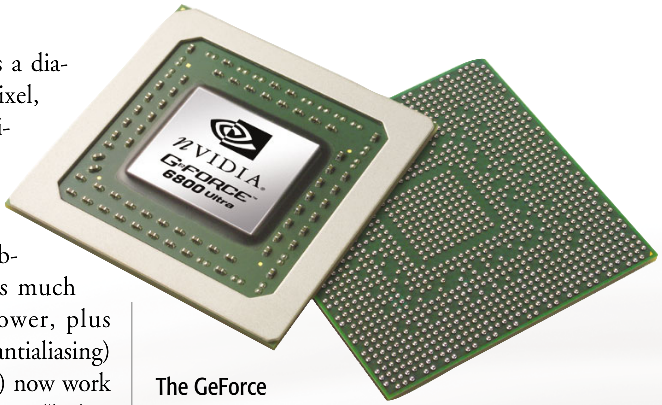

The GeForce 6800 Ultra.

Video

ATI had previously dominated the field with its VPE (video processing engine), going as far as implementing programmable shader functions for video streams. For NV40, it has packaged an integrated, programmable, on-chip video processor for handling video acceleration functions and video hardware encoding/decoding.

The VP comprises address, scalar, vector, and branch units (the vector unit is a 16-way SIMD vector unit), sporting its own separate pipeline system supporting dynamic branching and integer processing. The VP incorporates an adaptive deinterlacing engine with improved filtering and scaling techniques with video-deblocking capabilities and an integrated TV encoder, which should help challenge ATI's All-in-Wonder. NV40 also comes as a proper HDTV solution, so you can output HDTV streams in 720p, 1080i, 480p, or CGMS formats. (Just buy a $50 antenna.) The PVR functions include hardware AV synchronization, hardware MPEG-1/2/4 encoding/decoding, and WMV9 decode acceleration with motion compensation and in-loop deblocking. The VP also provides a nifty inverse telecine 3:2 pull-down screen for watching telly on the PC, color space conversion, gamma correction, noise reduction, frame rate conversion, and motion-adaptive deinterlacing, all in hardware.

Halo: Combat Evolved

In the past, this DX8 class title suffered from rather poor performance, even on high-end 3D cards. It uses PS1.1/1.4/2.0 (we tested in 2.0), plus Vertex Shader 1.1. Previously, R3xx- and NV3x-class cards would show frame rates (usually 60fps or below) extremely close to each other at all resolutions below 1,600 x 1,200 x 32. The GeForce 6800 throws a spanner in the works here; the game actually becomes less of a jerkovision experience, with extremely significant boosts in frame rates at 1,280 x 1,024 and 1,600 x 1,200. (FSAA isn't supported, so we ran benchmarks in stock non-FSAA or AF modes.)

Halo: Combat Evolved (Pixel Shader 2.0)	**Radeon 9800XT**	**GeForce FX 5950**	**GeForce 6800**
1,024 x 768	60.5	55.7	92.5
1,600 x 1,200	28.9	27.7	61.6

Gun Metal

Yeti Studios' publicly available benchmark demo, based on its Gun Metal title, was created primarily to be extremely GPU dependant. It's a mixed bag of tests with its DX8-class Pixel Shader 1.1 and DX9 Vertex Shader 2.0. It was designed to bring GPUs to their knees. Thus, the low frame rates obtained aren't intended to offer real-world performance of this gain. The 6800 Ultra has a strong lead, especially at the higher resolutions with AF and AA turned on.

GunMetal Benchmark 1	**Radeon 9800XT**	**GeForce FX 5950**	**GeForce 6800**
1,024 x 768 4XAA/8XAF	58.8	39.3	76.7
1,280 x 1,024 4XAA/8XAF	43.5	32.1	70.7
1,600 x 1,200 4XAA/8XAF	29.8	26.2	54.4
GunMetal Benchmark 2	**Radeon 9800XT**	**GeForce FX 5950**	**GeForce 6800**
1,024 x 768 4XAA/8XAF	64.5	48.7	73.2
1,280 x 1,024 4XAA/8XAF	62.7	42.1	74.9
1,600 x 1,200 4XAA/8XAF	45.8	35.1	66.7

The Card

A standard 6800 at $299 is also on the cards and will sport 12 pixel pipelines, but that will come later. Our 6800 Ultra reference design board was one power-hungry beast, requiring at least a 480W power supply. (We tried it with a 450W PSU, and everything was hunky-dory.) It also requires plugging in two molex connecters, and it is not small form factor friendly. Power usage

Return to Castle Wolfenstein: Enemy Territory

Enemy Territory is a free add-on to Return to Castle Wolfenstein and has become a popular online multiplayer. Even if its 3D engine is vintage, it has been modified, so it taxes the GPU more than the original. Hence, it's more reflective of today's games based upon id Software's engine. The first four of the five test modes used on the GeForce 6800 almost went unnoticed in terms of a performance penalty. Only at 1,600 x 1,200 with 4XAA and 8XAF did the performance drop significantly, but its still miles better than the next best score.

Return to Castle Wolfenstein: Enemy Territory	Radeon 9800XT	GeForce FX 5950	GeForce 6800
1,024 x 768	99.1	111.5	109.4
1,600 x 1,200	90.9	91.4	109
1,024 x 768 4XAA/8XAF	78.6	79.8	108.8
1,280 x 1,024 4XAA/8XAF	56.3	61.9	100.8
1,600 x 1,200 4XAA/8XAF	40.3	38.2	70.9

Serious Sam: TSE

This OpenGL game is beginning to show its age when compared to Far Cry. It still remains a valid OpenGL test, however. Once again, the GeForce 6800's 16 pipes and extra bandwidth take the lead, and the card barely drops from 97.8fps at 1,280 x 1,024 with 4XAA and 8XAF down to 75.9fps at 1,600 x 1,200 with the same settings.

Serious Sam: The Second Encounter	Radeon 9800XT	GeForce FX 5950	GeForce 6800
1,024 x 768	101.7	121.6	126.1
1,600 x 1,200	65.2	83.5	119.4
1,024 x 768 4XAA/8XAF	71.3	82.4	117.4
1,280 x 1,024 4XAA/8XAF	51.2	59	97.8
1,600 x 1,200 4XAA/8XAF	38.3	43.4	75.9

Call of Duty

Yet another 3D FPS based on a heavily modified version of the Quake 3 engine, CoD was many a gamer's 2003 game of the year. The GeForce 6800 enjoys a lead, but owning one doesn't buy you that much until you turn on 8XAF and 4XAA. You could argue that no one cares about or needs FSAA in a fast-paced action shooter where you don't have the time to check for jaggies. But the scenery does look a little cleaner and crisper.

Call Of Duty	Radeon 9800XT	GeForce FX 5950	GeForce 6800
1,024 x 768	141.9	163.2	164.5
1,600 x 1,200	127.2	145.7	157.5
1,024 x 768 4XAA/8XAF	140.1	147.5	165.3
1,280 x 1,024 4XAA/8XAF	116.6	117.5	162.8
1,600 x 1,200 4XAA/8XAF	88.5	91.7	151.3

is up from NV38's peak 80W to 110W for the Ultra board.

With PCI Express still somewhat on the horizon (sort of like that French ship in "Master and Commander"), AGP 8X is still the order of the day, me hearties. The 6800 Ultra will take up the PCI slot closest to your AGP slot. The cooling solution looks elaborate with an improved heat pipe cooling solution that's actually relatively quiet. (No Dustbuster here.) The heatsink didn't inflict any heat-related injuries and was fairly cool to the touch. Overclocking proved largely unsuccessful on our card, but we can hope for further improvements from OEMs. Our board came with dual DVI connectors, which hopefully is the shape of things to come at the high end.

Theoretical Fill Rate & Pixel Shader Performance

It's one thing acknowledging specs on paper and another putting them into practice, which is what we did using 3DMark03's specific fill rate and pixel/vertex shader isolation tests. The theoretical maximum fill rate for the 6800 Ultra is listed at 6.4Gtps. When pushed hard in specific single and multitexturing tests, it hit 2848.5Mtps for single texturing and 6177.2Mtps for multitexturing. Compare this to the older 5950 Ultra, and the 6800 Ultra looks like its 16 pipes make it a multitexturing monster. The battle between the previous generation cards for better pixel and vertex shading was up for debate, but that's no longer the case with the 6800 Ultra, which has twice the vertex shader performance as the 5950 and comes up trumps by 2X in the vertex shading test and more than 2X in the PS2.0 tests. This improvement in vertex shading for NVIDIA is drastic.

General Performance Summary

After considerable time benchmarking and playing games on the 6800 Ultra, it's apparent the card barely breaks a sweat going from 1,280 x 1,024 to 1,600 x 1,200. If you don't have a 21-inch monitor that supports that high of a resolution, you're missing out. Furthermore, the performance hit associated with turning on 8XAF (in some cases brilinear, groan) and 4XFSAA is anywhere from extremely slight to still acceptable and well worth using in all the games we tested. One can't say that of the previous generation of hardware.

The GeForce FX 5950 Ultra and Radeon 9800XT seem to hit a fill rate wall when jumping to 1,600 x 1,200 in some games. Turning on better filtering with 8XAF and then attacking those jaggies with 4XAA puts stress on the previous generation of cards, especially at 1,280 x 1,024 and above. We found it wasn't worth the hit in quite a few titles over the past few months, especially FPS, which incur excessive slowdowns during more intense action. With the NV40 Ultra's higher fill rate and pipes, this was no longer true. When moving from less complex to more complex scenes—from indoors to outdoors in Far Cry, for example—the 6800 Ultra was far more seamless in terms of slowdowns. Not only is the 6800 Ultra the clear winner by a significant margin, it also shows that it paid off for NVIDIA to address pixel shader performance, which reflects the largest deltas of performance in PS1.1/1.4/2.0 games. The improved algorithms help to make it an AA and AF beast at the higher resolutions, and it's numbing to see these titles become CPU bound.

In most cases 60fps is considered optimal when playing a game at a resolution and visual quality setting of your choice. When looking at our benchmarks, consider all scores below 60fps as a FAIL and those above 60fps as a PASS. The NV40 becomes even more compelling at that point. Do the same with a 30fps ceiling, which we consider the absolute minimum acceptable frame rate, and even then some FAILS show up. If 1,600 x 1,200 8XAF and 4XAA are important to you, as they should be at $499, your choice becomes fairly clear, even if you already own a Radeon 9800XT or GeForce 5950 Ultra—although not crystal clear until the next generation of ATI hardware shows up. (For more on ATI's X800, see page 20.)

NVIDIA has come back very strong with NV40, and it's going to take a strong effort from ATI to top this. Justifying a $499 price tag for a 3D accelerator isn't easy, but NVIDIA has done an incredible job doing so with the GeForce 6800 Ultra.

(For more benchmark scores, plus our "The Tale Of The Specs" chart, subscribers can go to www.cpumag.com/cpujun06/NV40.)

---

by Alex "Sharky" Ross

Far Cry

This next-gen 3D game separates the men from the boys; it's undoubtedly the game of the year, thus far. It sports the snazziest 3D engine to date in any 3D FPS. Making full use of PS2.0 and Vertex Shader 2.0 code paths, it supports DOT3 bump mapping, Projective EMBM, shadow maps, stencil shadows, and more. We created our own time demo on the Research map, which starts in the jungle and works back indoors to simulate a typical few minutes of gameplay. Far Cry hit the 9800XT and 5950 pretty hard as soon as AF and AA were turned on. So hard that 1,600 x 1,200 with FSAA and 8XAF enabled was unplayable at 15.4fps. The penalty hit on the 6800 wasn't nearly as rough. One point regarding image quality concerns the water, which looked superior on the 9800XT.

Far Cry: Research Level (indoor/outdoor)	Radeon 9800XT	GeForce FX 5950	GeForce 6800
1,024 x 768	68.2	55.5	92.4
1,600 x 1,200	38.9	29	84.2
1,024 x 768 4XAA/8XAF	46.4	33.6	82.3
1,280 x 1,024 4XAA/8XAF	30.2	23.9	66.7
1,600 x 1,200 4XAA/8XAF	21.9	15.4	49.8

MS Flight Simulator 2004

A rather demanding 3D engine but one that's renowned for being CPU bound, although there does seem to be a delta between the 6800 Ultra and the rest once FSAA and AF are turned on.

Flight Simulator 2004	Radeon 9800XT	GeForce FX 5950	GeForce 6800
1,024 x 768	55.1	58.4	58
1,600 x 1,200	53.4	57.3	56
1,024 x 768 4XAA/8XAF	50	50.6	57.3
1,280 x 1,024 4XAA/8XAF	39.6	39	56.6
1,600 x 1,200 4XAA/8XAF	31.8	30.8	54.1

X2 - The Threat

This DX8 title utilizes a rather complex 3D engine with particle effects and actually utilizes stencil shadow techniques that put UltraShadow technology to good use. The contest wasn't even close at any settings we tested.

X2: The Threat	Radeon 9800XT	GeForce FX 5950	GeForce 6800
1,024 x 768	67.6	71.7	95.6
1,600 x 1,200	41.5	48	66.4
1,024 x 768 4XAA/8XAF	43.9	49.8	77.6
1,280 x 1,024 4XAA/8XAF	32	35.5	60.4
1,600 x 1,200 4XAA/8XAF	24.7	27.4	49

Super wireless performance– it runs in our family.

Faster access to multi-player gaming
Video streaming without jitters
Effortless wireless connection – gets anyone up and surfing quickly
Superior wireless range

Total 802.11 Super G™ ATHEROS

Super G™ Technology – 10X Faster

108 Mbps 2.4 GHz 802.11 g

WGT624
108 Mbps
Firewall Router

WG511T
108 Mbps
PC Card

WAG511
108 Mbps
Dual Band
PC Card

WAG311
108 Mbps
Dual Band
PCI Adapter

Introduce your family to our NETGEAR® 802.11g family that offers exceptional range, compatibility, security and super speed.

Now you can choose 802.11g technology with confidence, thanks to NETGEAR. Recipient of two CNET awards for exceptional performance and ease of use, our 802.11g family runs circles around other companies' wireless networking products. It's fully compliant with the recently ratified 802.11g standard and compatible with all 802.11b and 802.11g products.

Our NETGEAR 802.11g family. The perfect addition to *your* family. For details, stop by for a visit at **www.netgear.com/go/g**.

54 Mbps 2.4 GHz 802.11 g

c|net Editors' Choice APR 2003

WG602
54 Mbps
Access Point

c|net Editors' Choice APR 2003

WG511
54 Mbps
PC Card

WG311
54 Mbps
PCI Adapter

WGR614
54 Mbps
Router

WGE101
54 Mbps
Ethernet Bridge

WG121
54 Mbps
USB Adapter

Everybody's connecting.™

NETGEAR®

©2003 NETGEAR, Inc. NETGEAR is a registered trademark and "Everybody's connecting." is a trademark of NETGEAR, Inc. All rights reserved. Other brands and products are trademarks of their respective holder/s. E134 Smart Computing

by Warren Ernst

The Bleeding Edge Of Software

Inside The World Of Betas

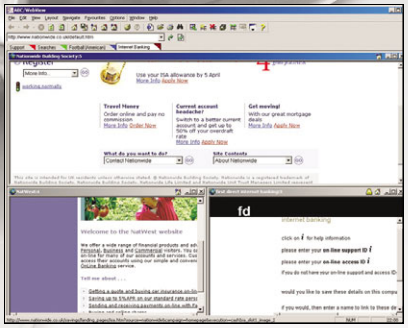

Official product name: ABC/WebView
Version # previewed: 4.1.3.7 Beta
Publisher: ABC Enterprise Systems
Developer and URL: ABC Enterprise Systems; www.abcwebview.com
ETA: Q2 2004
Why you should care: If you power-browse the Web, this may be the tool for you.

ABC/WebView 4.1.3.7 Beta

Just when you thought the browser had been innovated out, along comes something new—sort of. Like MyIE2 and others, ABC/WebView uses the core Internet Explorer rendering engine to decode HTML and display Web pages. In this sense, there's nothing new. In other ways, ABC/WV displays multiple Web pages in a new way, benefiting power users.

Like Mozilla, Opera, and Safari, ABC/WV uses tabs to display multiple pages within the same parent browser window. Unlike those browsers, each ABC/WV tab can also display several pages, arranged in tiles, rows, or columns. So you can have one tab called News and have three child windows in it displaying CNN, MSNBC, and The Onion. You can also have a second tab called Finance and have windows displaying Citibank and Morning Star.

ABC/WV calls these sets of tabs and child windows "Layouts." The app can save and reload Layouts through menu commands or when you start the program. It can also keep track of where you've scrolled on a Web page so that reloads always appear at the same place on a page. Additionally, ABC/WV allocates more bandwidth to the pages in the topmost view, meaning you can still browse in the foreground while background pages load, even over a dial-up connection.

Other features you'd expect in a modern browser are here, including pop-up window suppression, timed auto reloading, and banner ad blocking. The toolbar buttons, however, are a throwback to Windows 95, and some are positively baffling. Overall, ABC/WV works well and is suited to those who browse many pages at once. ▲

JukeBox Synchronizer 3.2.6 Beta

Many portable MP3 players look like a drive letter to Windows when they're plugged into your PC. If yours does, maybe you've wondered if there's a better way than just dragging and dropping to synchronize your PC-based music library with your portable library. After all, dragging and dropping doesn't help with playlist creation and randomizing. If so, JukeBox Synchronizer will thrill you.

JBS looks at the files in a designated folder or set of folders on your PC and compares them to files in another folder or set usually located on your portable device. It also works with a network drive. JBS then copies new or updated files in either or both directions, depending on your preferences.

What separates JBS from other tools is that it's geared for music. If your portable player's collection is smaller than your library, JBS will randomly select files that will fit and copy them to the player, offering a different slice of your collection on your player each time. It can also create and update playlists of your entire library or the portion that gets copied over. If your entire collection fits, the app does a very good job of randomizing the playlists.

This beta adds a working playlist editor (fully drag and drop) and can change pathnames from absolute (C:\80s\song.mp3) to relative (\80s\song.mp3) so you can synchronize playlists, too. The app is easily worth $25. Fortunately, it's free. ▲

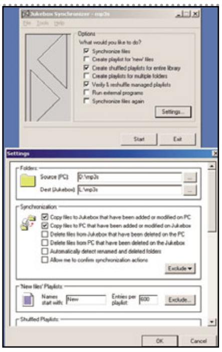

Options
What would you like to do?
☑ Synchronize files
☐ Create playlist for 'new' files
☑ Create shuffled playlists for entire library
☐ Create playlists for multiple folders
☑ Verify & reshuffle managed playlists
☐ Run external programs
☐ Synchronize files again

Folders
Source (PC): D:\mp3s
Dest (Jukebox): L:\mp3s

Synchronization
☑ Copy files to Jukebox that have been added or modified on PC
☑ Copy files to PC that have been added or modified on Jukebox
☐ Delete files from Jukebox that have been deleted on the PC
☐ Delete files from PC that have been deleted on the Jukebox
☐ Automatically detect renamed and deleted folders
☐ Allow me to confirm synchronization actions

'New files' Playlists
Names start with: New
Entries per playlist: 600

Official product name: JukeBox Synchronizer
Version # previewed: 3.2.6 Beta
Publisher: C-Flash Software
Developer and URL: C-Flash Software; www.cflashsoft.com
ETA: Q3 2004
Why you should care: Maybe the best way to synchronize your music library to a portable player.

MediaMonkey 2.2.736 Beta 4

Looking for more in a media player but don't want to pay a lot (or anything)? Check out MediaMonkey, a Windows-native, iTunes-like player. It includes auto categorizing of your tagged MP3 library, auto playlist generation, most of the features that make iTunes great, plus many extras.

Once MM scans your hard drive for supported media, it sorts your music into categories based on information gleaned from your ID3 tags. So, you can see all of your Phil Collins music in one neat list, regardless of your actual hard drive folders, file names, or album sorting. Categories include Artist, Album, Genre, Year, Title, Classification, and Your Rating (from one to five stars). MM also cleans up your tags and uses the same categorization scheme to locate files with unknown artists/titles/etc, duplicate files, bad files, and more. Although MM's actual tag-editing tools are slightly primitive compared to dedicated file retaggers, its ability to locate problem files may be unmatched.

The app also plays well with others. Although it has a built-in media player (which looks so much like Winamp 2 you can skin it using Winamp 2 skins), MM can directly hook into Winamp, sending files and playlists directly to that player. It also hooks directly into Nero software, making it easy to make an MP3 music CD or rip a CD and tag the music directly from Gracenote's CDDB online database. MM can synchronize your libraries on a PC and portable device, and there's a file-format converter that works extremely well. There are also visualizations, full skin support, about 20 external plug-ins, and much more.

My beta was very solid in terms of crashes and memory leaks, but there were some little niggles. For example, when playing shuffled music, MM didn't remember what the last song was when I clicked Previous Song. Still, MM is the most innovative and feature-filled player I've seen yet for Windows, and its free price is simply amazing. ▲

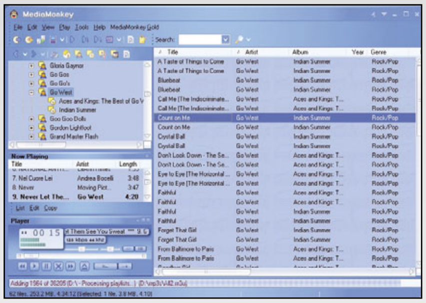

Official product name: MediaMonkey
Version # previewed: 2.2.736 Beta 4
Publisher: Ventis Media
Developer and URL: Ventis Media; www.mediamonkey.com
ETA: Q2 2004
Why you should care: At last, an iTunes-like program with more features and less bloat than iTunes.

GX::Transcorder 1.82.1274 Beta

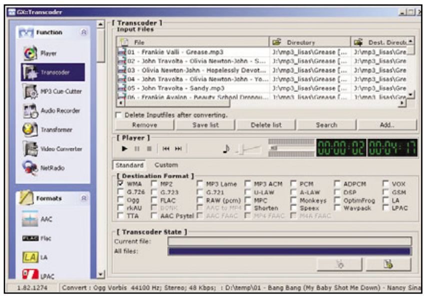

Official product name: GX::Transcorder
Version # previewed: 1.82.1274 Beta
Publisher: GermaniXSoft
Developer and URL: GermaniXSoft; www.board-24.de
ETA: Q2 2004
Why you should care: The ultimate free media-file-format converter.

When free software does a great job at tough tasks, it's easy to overlook a tough installation and somewhat-bizarre interface sprinkled with English and German. Such is the case with GX::Transcorder, a multimedia file-format converter specializing in AV files.

The app's large installation includes many codecs and two Windows Media file libraries (one in the Windows Media SDK). However, not all the codecs you'll want are included, and the Microsoft-sourced files that claim to install on Windows 98/Me won't, making GXT a Win2000/XP product only. You can get the extra codecs by reading the readme files and error messages, but the pages you're directed to don't make it clear what exactly you need to download or where to place the downloaded files.

GXT does offer a lot. Properly equipped, it can convert around 40 audio formats, from MP3 and Ogg Vorbis to the rarer G.721 and M4A FAAC. Where permitted, internal data tags are maintained. GTX can convert AVI to MPEG to WMV in any direction, allowing for a few dozen subvariants of each. DVD, VCD, and SVCD extensions are available for burning discs, but you'll need another app to write.

The package also includes a universal audio recorder, which records anything passing through the audio card to a WAV, MP3, or other type of file. There's also an Internet radio player (with hardly any stations worth listening to), an audio special effects generator/filter, and basic media player. These work, but there's nothing special about them.

Conversions were first rate and worked on several "unconvertible" files in my library. Given the app's price, I can overlook its problems and recommend it. ▲

Send Us Your Betas: Know of software in the beta stage that's deserving of some attention?
Let us know. We'll take a look at it and possibly give it a go-round. Send your prospects to bleedingedge@cpumag.com.

UP TO SPEED

Upgrades That'll Keep You Humming Along

This month sees some significant upgrades, including for the popular Eudora and open-source OpenOffice, plus updated driver releases from Creative, Intel, NVIDIA, and NetGear.

Updates

Audiograbber 1.83

This tool for ripping, encoding, and normalizing audio from CDs goes freeware and gets some bug fixes.

Get it at: www.audiograbber.com-us.net

Dreamweaver MX 2004 7.0.1

This update fixes bugs, restores the Timeline feature, and improves stability.

Get it at: www.macromedia.com

Driver Cleaner 3.0

This utility cleans hardware drivers from your hard drive and the Windows Registry, where left-behind items and OS instructions can conflict with newer drivers and other hardware. This version improves removal of ATI and NVIDIA graphics drivers and Creative sound card installations.

Get it at: www.driverheaven.net

Eudora 6.1

Eudora gets spam filters, an integrated Google search box, better searching into message threads, email filing based on keywords, and improved importing from Outlook. A free ad-supported version (with fewer features) is available.

Get it at: www.eudora.com

KlipFolio 2.5

This free memory-resident information gatherer lets you subscribe to thousands of news feeds and blogs, pushing them to the desktop in a pop-up reader. This version includes audio news alerts, scroll bars in the news feeds, and customized fonts.

Get it at: www.serence.com

Info Select 8

Better word processing, image editing, resizing, and importing are among the scores of changes in this major version upgrade for this information manager, organizer, and free-form database.

Get it at: www.miclog.com

Internet Download Manager 4.00

This tool for accelerating, recovering, and resuming downloads gets a rebuilt engine, more speed, new browser integration, and usability improvements.

Get it at: www.internetdownloadmanager.com

OpenOffice.org 1.1.1

This free, open-source suite includes most of the modules and functionality of Microsoft Office. The update fixes minor and rare bugs and adds new features.

Get it at: www.openoffice.org

SYSmark 2004 Patch-1

This first major patch for the venerable system benchmark fixes problems often encountered when it runs its complex scripts. Improved hardware detection and upgraded Help files are also included.

Get it at: www.bapco.com

Visual Basic 6.0, Visual C++ 6.0 With Visual Source Safe 6.0d Service Pack 6

Recommended by MS to bring users of these apps up-to-date with the latest code, this service pack contains many bug fixes.

Get it at: www.microsoft.com

Winamp 5.03a

This refresh for the recently released Winamp 5 fixes a bug in the track repeat feature and problems with applying free-form player skins.

Get it at: www.winamp.com

Driver Bay

Audigy 2/ZS EAX 4.0 Advanced HD Driver Patch

For Audigy 2 and Audigy 2 ZS models only, the patch resolves problems with THX Setup Console, synchronizes EAX Advanced HD settings and the Creative Graphics Equalizer, and enhances performance of games running under OpenAL API extensions. A separate EAX 4.0 patch is available for original Audigy cards.

Get it at: www.creative.com

Intel Extreme Graphics 2 14.1

These display drivers are for integrated mobo graphics using the Intel 865G, 852/855 GM/GME, 845G, and 830MG chipsets. This version resolves minor and rare occurrences often involving graphics resuming after Standby and Hibernation modes or when running on dual displays. Many bug fixes are also included.

Get it at: www.intel.com

NetGear WG511 802.11g Adapter Software 2.1.0

For NetGear's popular 802.11g PC Card, the drivers comply with official Wi-Fi Protected Access standards and the final IEEE 802.11g standard.

Get it at: www.netgear.com

NVIDIA ForceWare 56.72

The latest WHQL Certified (except on some earlier models) GeForce/TNT graphics card drivers support GeForce 5700/Ultra models. The drivers include DirectX9 and OpenGL 1.5 support and upgrade HDTV output from cards that support it.

Get it at: www.nvidia.com

NVIDIA System Utility 1.08.5

For nForce2 and nForce3 mobos, the utility tweaks and monitors system voltage, temps, overclocking, and more from within Windows rather than the BIOS.

Get it at: www.nvidia.com

by Steve Smith

Enlisting A Video Encoder

For When It's Time To Compress & Decompress

Video encoders can compress a video file and save it in a variety of file formats and sizes while retaining most of the visual quality of the original file's images. Compressing video is crucial because just 10 seconds of ordinary NTSC video can gobble up to 300MB of storage space. If you want to distribute video on a CD/DVD, as a streaming video file, or as an email attachment, you need to compress it.

For maximum flexibility an encoder must accept video from various sources, such as a videocassette or digital video camera. You must also be able to save that video file in different file formats. For example, you may need one format to save the video as a streaming video file but a different format to save it to DVD. To decide which video encoder is best for your needs, there are some common and special features you should know about.

How Video Compression Works

To compress video, you need a video codec, which is an algorithm that specifies how to compress video for storing and how to decompress it for playback. The efficiency of a codec determines how small it can compress the video without losing its original visual appearance. Compressing video involves a trade-off between small file sizes and image quality. The smaller the file, the lower the image quality. Most video codecs use lossy compression, which means removing part of your video. The less video a codec stores, the smaller the file.

A codec can't indiscriminately chop out chunks of your video without completely destroying the video image, so codecs selectively remove parts by comparing the video frame by frame. At the simplest level, the codec only saves images that change from one frame to the next. Because a background may stay the same throughout multiple frames, the codec only stores the background from a single frame, cutting the size of the video considerably.

Video Encoder Glossary

AVI (Audio-Video Interleaved)–A video file format commonly used by Windows computers.

bit rate–The amount of space needed to store (or send) one second of information, usually measured in kilobits per second.

codec–Short for compression/decompression, technology that compresses and plays back video.

CBR (constant bit rate)–A method of encoding video content that's independent of the visual images.

key frame–A frame that a codec uses to save images that don't change, such as a background.

MPEG (Motion Picture Experts Group)–A video compression standard used to store video. The MPEG file format includes the MPEG-1, MPEG-2, and MPEG-4 standards.

VBR (variable bit rate)–A method of encoding video content based on the complexity of the visual images. VBR is usually slower than CBR and may require two passes to provide optimum visual quality.

Codecs can also cut the file size by reducing the number of colors saved, altering the video's frame rate, and lowering audio quality. These methods involve cutting information to save storage space, but if a codec is efficient enough it can reduce colors so you won't notice or lower audio quality to eliminate sounds, you can't hear anyway.

Give & Take

No matter the codec you use, there's always a trade-off between compression level, video quality, and compression/decompression speed. The compression level defines how much a codec can compress a file. Video quality determines how much visual information the compressed file retains from the original. Compression/decompression speed measures how fast the codec can compress the video and decompress it (to play it back).

If you compress too much, video quality suffers. If the codec is too slow in analyzing video, the compression/decompression speed can take too long. Some users might tolerate a slow codec for compression, but no one will wait several hours for the codec to decompress the video to watch it. If a codec takes equal time to compress and decompress a video, it's said to be symmetrical. If there's a big difference between the two times, the codec is said to be asymmetrical.

Although there are many different codecs, no single codec is best for all types of video. Some are better suited for outdoor video images. Others are better at distributing video on CD/DVD. And still others are better for creating streaming video. For maximum flexibility most video encoders offer several codecs so you can choose the best one based on the task.

Tweaking A Codec

A codec alone can't do anything without a video encoder program. The codec takes care of compressing the video while the video encoder program lets you modify the way the codec works. Most video encoders can guide you in compressing a video without knowing anything about codecs, but modifying codec settings can let you wring the last bit of compression out of your video. Some settings include bit rate (or data rate), frame rate, frame size, and key frame.

The bit rate measures how much space (usually in kilobits) is used to store one second of video. A high bit rate improves video quality but increases the file size. Depending on how you plan to distribute your video, there may be a maximum bit rate value. For example, a VCD can only playback video at a rate of 1,100Kbps, so storing a video that specifies a higher bit rate than 1,100 is pointless.

Many encoders use CBR (constant bit rate) compression or VBR (variable bit rate) compression. CBR reduces the bit rate uniformly throughout the entire video. Although this means the codec can compress video quickly, it also means the codec treats simple images (the moon against a black background) the same as more complicated images (a forest on fire with airplanes flying overhead), so the video quality of the more complicated scenes may suffer.

VBR alters the bit rate depending on the complexity of the video image. This can improve video quality but means slower compression time. For even higher video quality, some encoder programs offer two-pass VBR, meaning the program analyzes your video file twice to get the maximum video quality and smallest file size possible, but at the expense of much slower compression speed.

The frame rate defines how many frames per second are used to store your video. The fewer frames used, the lower the quality and smaller the file. Because many computers can't play video faster than 16fps, a lower frame rate can preserve quality while reducing the file.

Frame size measures the physical dimensions that your video appears in. A large frame size takes up more storage space. Key frame defines how many frames the codec waits before capturing another unchanged image, such as a background. Lower the key frame setting, and the codec captures unchanging images more often to improve video quality, but with a greater file size.

When choosing a video encoder, don't be fooled by the minimum requirements. You could use Windows 98 with a Pentium III processor and 128MB of RAM, but a P4 with 256MB of RAM really should be the minimum requirements because video files gobble up RAM, hard drive space, and processing power. We tested three video encoders geared for a wide range of users using a 2.4GHz P4 system with 512MB of RAM and 1.8GHz AthlonXP system with 512MB of RAM. Both computers ran Windows XP. Obviously, a computer with a faster processor or more RAM would likely run through the test files faster, but these systems represent an average user's typical configuration.

Pegasys TMPGEnc 2.5

To learn how a video encoder can compress and convert video files, grab the free trial version of TMPGEnc (Tsunami MPEGEncoder). While most commercial video encoders let you compress video into a variety of file formats, TMPGEnc only lets you compress video files into MPEG or AVI format.

Despite its spartan user interface (only File, Option, and Help menus), TMPGEnc provides a surprising amount of features for compressing and converting video files. To speed up compression, TMPGEnc can use MMX (Multimedia Extension)-enabled Intel processors, plus support multiple CPUs.

To further reduce the size of your video files, TMPGEnc provides a handful of filters that can clip the edges of a video's frames, filter out excessive sounds, and correct color tone. Although TMPGEnc's commands are few and relatively easy to navigate, it still offers a helpful Project Wizard to help users compress files. If you need to compress multiple videos, TMPGEnc even offers a special batch mode. This lets you just define your compression settings, load the video files you want to compress, and leave your computer running overnight.

At just $88 for a boxed version that comes with a printed manual (or $48 for a download version), TMPGEnc is a perfect program for hobbyists learning about video compression. Best of all the free 30-day trial version lets you experiment with your own video files to see how different settings can change the quality and size of your files. If you only need to work with MPEG or AVI file formats, TMPGEnc may be all you need. For more than this, you'll have to look at other programs. The app runs on Windows 98/Me/NT/2000/XP.

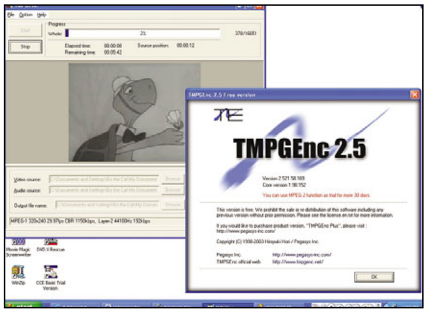

TMPGEnc 2.5
$88 (boxed); $48 (download)
Pegasys
www.pegasys-inc.com

Sorenson Media Sorenson Squeeze 3.5 Compression Suite

Sorenson Squeeze packs enough features to handle all the video compression needs for the casual hobbyist and professional. The app can save your video in practically any file format, including QuickTime, Windows Media, Real Media, MPEG, and even Macromedia Flash. Besides supporting various file formats for output, it can also accept video stored in AIF, ASF, AVI, DV, MPEG, QuickTime, and Windows Media formats.

To optimize video size and quality, Sorenson Squeeze can automatically choose the best settings based on how you plan to distribute and play your video. For example, the app provides default settings for creating a streaming video suitable for both dial-up and broadband users. It can also output several different file types simultaneously, such as converting an AVI file to an MPEG and QuickTime file.

Although the program lacks a step-by-step wizard, the UI is simple enough to choose the target distribution for your video (such as a specific file format or a media), plus provide default settings optimized for your target file format or distribution medium. Naturally you can alter these settings. When you change

Testing Video Encoders

Benchmarks measuring file size may look nice, but they can vary drastically depending on the types of files you use. Still, the results we saw can give you a rough idea how different programs compare and the results you might expect with your own video files. The four test files we used include a 1.6MB home movie captured with a Sony digital video camera, a 29.6MB AVI file, a 140MB MPEG-1 video, and an 11.1MB QuickTime movie. In general, converting video from one file format to an MPEG will result in a larger file, while compressing a file originally stored as an MPEG can result in a smaller, compressed version. ▲

	Pegasys TMPGEnc 2.5	Sorenson Squeeze 3.5 Compression Suite	Cinema Craft Encoder Basic
Converting a 1.6MB digital video movie to an MPEG-1 file resulted in a:	14.6MB file	16.7MB file	Unable to load
Converting a 29.6MB AVI file to MPEG-1 resulted in a:	90MB file	81.9MB file	139MB
Converting a 11.1MB QuickTime file to an MPEG-4 file resulted in a:	Unable to load	28.9MB file	Unable to load
Compressing a 140MB MPEG-1 file resulted in a:	137MB file	71.7MB file	Unable to load

compression settings, you alter the final file size. But if you want to store your video on a CD that only has a fixed amount of space remaining, you might need to fiddle with settings until you get your video to a size that fits.

Adjusting compression settings can be time-consuming, so Sorenson Squeeze lets you specify a maximum file size for your compressed video. The program then chooses the best compression settings to smash your video down to that file size. The app's hefty $499 price puts it out of the casual enthusiast's range, but its cross-platform compatibility (Win98/Me/2000/XP and Mac OS X) and intuitive UI make it an indispensable tool for anyone who needs to distribute video.

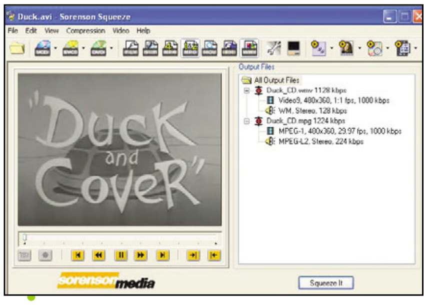

Output Files
- All Output Files
  - Duck_CD.wmv 1128 kbps
    - Video9, 480x360, 1:1 fps, 1000 kbps
    - WM, Stereo, 128 kbps
  - Duck_CD.mpg 1224 kbps
    - MPEG-1, 480x360, 29.97 fps, 1000 kbps
    - MPEG-L2, Stereo, 224 kbps

Sorenson Squeeze 3.5 Compression Suite

$449
Sorenson Media
www.sorenson.com

Cinema Craft Encoder Basic

Cinema Craft Encoder Basic ($59) is a stripped-down version of Cinema Craft's professional video encoder ($1,950). You can download a free trial version of Encoder Basic, but if the trial version is anything like the commercial version, you may want to spend your money on a different video encoder.

The trial version claims to accept AVI and QuickTime files, along with digital video. While loading AVI files worked fine in the trial version, loading both QuickTime files and digital video constantly displayed an error message.

Encoder Basic encoded AVI files quickly and easily. The program's simplified UI involves loading a file and clicking an Encode button to convert it to an MPEG, which is the only output file format supported. While the UI is fine for those who just want to convert a video file into MPEG format, it prevents users from tweaking the compression settings to optimize the video compression. Worse, the program's MPEG version of AVI files consistently displayed odd color bands across the screen, along with jerky video with chunky images. The application even mangled parts of the audio soundtrack.

With limited file-format support and unacceptable video and audio quality, Encoder Basic may come from well-known and respected pedigree, but it ultimately failed to provide even minimally acceptable videos. The program works with Win98/Me/NT/2000/XP.

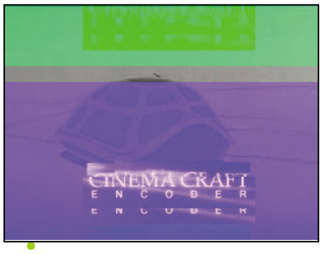

Cinema Craft Encoder Basic

$58
Cinema Craft
www.cinemacraft.com

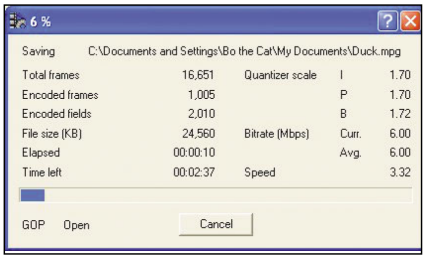

```
Saving          C:\Documents and Settings\Bo the Cat\My Documents\Duck.mpg
Total frames         16,651    Quantizer scale   I     1.70
Encoded frames        1,005                      P     1.70
Encoded fields        2,010                      B     1.72
File size (KB)       24,560    Bitrate (Mbps)    Curr. 6.00
Elapsed            00:00:10                      Avg.  6.00
Time left          00:02:37    Speed                   3.32
GOP   Open
```

What Do You Need?

More money definitely buys better performance. Still, unless you are a professional or an avid enthusiast, it is hard to justify buying a video encoder that can cost as much as a new computer. Fortunately a lower price doesn't necessarily translate into inferior quality but simply fewer features, which you may not need anyway. No matter your skill or budget, you should be able to find a video encoder that fits within your price range while offering the features you need. ▲

by Wallace Wang

CPU Ranking: 0 = Absolutely Worthless 2.5 = Absolutely Average 5 = Absolutely Perfect

SnapStream Media Beyond TV3

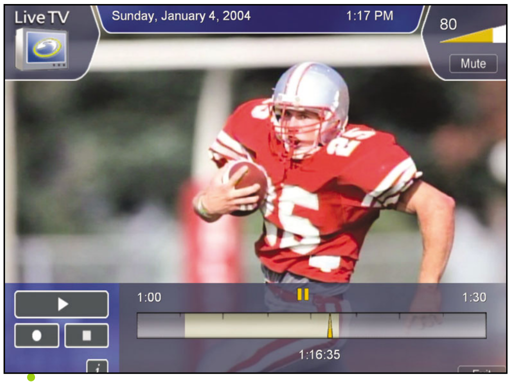

Beyond TV3
$69.99; $59.99 (download)
SnapStream Media
(713) 644-6240
www.snapstream.com
CPU CPU CPU CPU

That's right, Beyond TV3 can stream recordings to any machine on your LAN or across the Web via the Web Admin interface.

You're probably familiar with the basics of PVR apps: It's TiVo in a PC. Beyond TV3, SnapStream's follow-up to Personal Video Station, is the first app we've tried that clearly shows superiority over a Windows Media Center Edition system. MCE has a more integrated interface and better show search capabilities, plus an MCE machine comes with the necessary hardware to access the complete spectrum of satellite and digital cable, not just the basic cable accessible through your tuner card. Because we lacked an IR Blaster or similar tool, we used a Streamzap remote and tested in the basic cable range. If you don't have a compatible remote, Beyond TV3 works fine with a mouse and keyboard.

The app installs quickly and guides you through several screens to identify your hardware. Our Motorola digital cable box wasn't on SnapStream's list, but a similar model worked fine. Beyond TV3's first screen offers four program modes/modules to choose from. The Full Screen and Windowed modules are self-explanatory.

A Log module is a running operations diary you'll probably never need, and Web Admin is a browser-based module in which you can watch recordings, manage your present/future recordings, and browse and search the Program Guide through SnapStream.Net. Search criteria only includes titles, actors, and show descriptions. Use these rather than doing a global search. Otherwise, a general search for "Harrison Ford" will yield hits such as Rex Harrison and the seemingly unrelated "Sabrina, the Teenage Witch."

Web Admin also provides access to your ShowSqueeze log and most configuration settings, including the specifics of your streaming server. That's right, Beyond TV3 can stream recordings to any machine on your LAN or across the Web via the Web Admin interface. Conversely, you can schedule recordings and manage content from any remote location. Just type the Beyond TV3 box's IP address into your browser and/or log into your account at SnapStream.Net.

The Cool Stuff

No, not watching live TV, recording, time shifting, or 30-second skips. That's old news. We mean ShowSqueeze, which is the app's transcoding engine. All shows are recorded in MPEG-2 (a WMV option is noted in the manual but wasn't implemented in our version), but you can instruct the software to either replace or make a copy of the show in an MPEG-2, Windows Media, or (with an optional plug-in) DivX version to the quality and size settings you specify. For example, we used the WMV Pocket PC ShowSqueeze preset to make a copy of "Monster House." This turned a roughly 3GB recording into 15MB. The video quality was terrible, but hey, that's an hour of TV in your pocket.

Another killer feature is SmartSkip, an analysis tool that detects commercial blocks in your shows. With SmartSkip enabled, a show's timeline displays yellow blocks where the commercials reside. When you arrive at the beginning of the first commercial, press the up arrow on your remote and you immediately skip to the end of the last commercial to resume your show. We tried this on about a dozen programs, and it worked perfectly every time. SmartSkip alone is worth the price of this app.

Beyond TV3's interface is remarkably easy and intuitive. At the Best settings, we observed very little quality loss in either WMV or DivX from the original MPEG-2 files. Throughout our review, we had few questions, fewer complaints, and plenty of "oh, cooool!" moments.

The Beyond TV 3 interface still needs more integration. You should be able to change codec bit rates on the fly when you start a ShowSqueeze, but you can't because those options are in a different area. There's too much redundancy between the various modules, and the whole program should run from one interface. SnapStream tells us this sort of integration is in development.

A CD/DVD export module and automatic commercial deletion would be great, but SnapStream can only come so close to upsetting the broadcast white collars. As it is, we feel fortunate that SnapStream delivers a product that stores top-quality content in three widely used file formats without any DRM restrictions (yet) and helps you enjoy it with unprecedented flexibility. ▲

by William Van Winkle

CPU Ranking: 0 = Absolutely Worthless CPU CPU C 2.5 = Absolutely Average CPU CPU CPU CPU CPU 5 = Absolutely Perfect

iLife '04

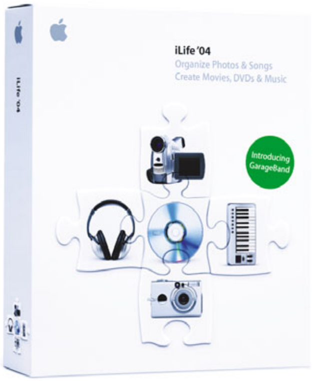

iLife '04
$49 (single computer);
$79 (five-computer Family Pack)
Apple Computer
www.apple.com/ilife

If you bought a Macintosh in the past few years, it came loaded with nifty Mac-only software, including iTunes, iPhoto, iMovie, and perhaps iDVD. Apple has updated those tools, added a new one, and rolled it all into a package called iLife '04.

Previously, three iLife components—iTunes, iPhoto, and iMovie—were available as free downloads. With the introduction of iLife '04, only iTunes is available free; you'll have to buy iLife (or a new Mac) to get the latest versions of the other programs. Free software is great, but it's hard to argue with iLife's $49 price, given the load of new features and speed improvements it brings.

GarageBand

GarageBand is the newest member of the iLife suite, and it's perhaps the one that initially will have many users scratching their heads. The app is a multitrack music editor that lets you create original songs using a library of MIDI instruments, 1,000 music loops, and 200 audio effects.

Depending on your perspective, this might sound incredibly cool or just useless. In the right hands, GarageBand's music sounds great (not like artificial computer-generated music), with realistic-sounding guitars, pianos, drums, and other instruments. You can make music by tinkering with the on-screen piano keyboard, but to really rock out, you will need a USB or MIDI musical keyboard, guitar, or other instrument.

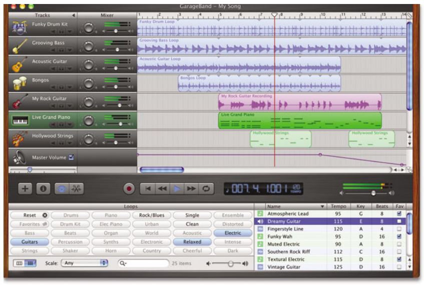

GarageBand is the newest member of the iLife suite.

It's hard to say where GarageBand falls in the spectrum between being a useful music tool and being a toy. A few glaring omissions keep it from being a true musician's tool. To start, it can't export MIDI tracks or show notes in standard musical notation. But it is a sophisticated tool for making music accessible even to people who don't know a sharp from a staff. If you can play an instrument or just have a song in your heart, GarageBand could be your gateway to composing it.

Apple also sells the GarageBand Jam Pack, a $99 add-on that adds 2,000 music loops, 100 instruments, and a gaggle of audio effects. At twice the cost of iLife itself, the Jam Pack is the domain of serious GarageBand enthusiasts.

iPhoto

iPhoto is a marvelous program for storing, cataloging, and manipulating digital pictures. New in this version is the ability to rate photos using a scale from one to five stars. You can then sort photos to see your favorites first. This version also adds Smart Albums, which automatically sort your photos based on their dates, titles, filenames, and other criteria (but unfortunately not EXIF data.) If you want to see the highest-rated photos you've shot with your Canon digicam in the last 90 days, for example, Smart Albums will show you.

In addition, there's a new sepia feature and a Rendezvous sharing feature that lets you view (but not edit) photos stored on other Macs on your network. The ability to rotate and delete pictures used to be irritatingly limited to the thumbnail view. Now you can rotate or remove photos in the full-screen view or even during a slideshow. Slideshows are also improved with new transitions, including a rotating cube and cool flipping mosaic effect, and you can use an iTunes playlist as an accompanying soundtrack.

New features aside, this new version is noticeably faster, especially with large sets of pictures. Apple says iPhoto can now handle up to 25,000 photos in an album. I can't vouch for that, but performance is much better when working with a 1,200-picture album on my 1.25GHz G4.

The rest of iPhoto remains largely unchanged, with plenty of options for printing photos, ordering prints, or creating custom picture books still intact. Like old baggage, some of the irritations of the previous version remain: The publish-to-Web feature is only compatible with the subscription-only .Mac service (that's OK, third-party utilities do a better job), and you can't import only selected images from your digital camera.

iMovie

For editing videos, there's iMovie. Like iPhoto, this new version delivers improved performance, especially when rendering transitions and other effects. A speedy machine helps; my 1.25GHz PowerMac kept me waiting while rendering, and effects could be

CPU Ranking: 0 = Absolutely Worthless 2.5 = Absolutely Average 5 = Absolutely Perfect

glitchy. Despite these minor problems, the machine was capable enough to make a movie without enormous frustration.

Another notable improvement is timeline editing with nondestructive edits. You can remove a section of video by simply sliding an arrow and then recover it by sliding the arrow back. The program also adds new effects and titles, an audio waveform editor, and alignment guides for meticulous synchronization of audio and video.

Once you complete your movie, the program can launch iDVD so you can burn it for posterity. You can also share movies by publishing it to your .Mac site (once again, it won't upload to other Web servers) or downloading it to a Bluetooth-compatible PDA or mobile phone.

iMovie can't compete with a professional video-editing tool, such as Adobe Premiere, but then again, iMovie is much easier to use. The app hits the mark for a consumer-level video editor; it's easy enough to learn but not so dumbed down that you'll outgrow it in a month.

iDVD

iDVD, the package's DVD-authoring application, is the logical companion to iMovie. You can create a movie and then burn it to DVD. The program's latest version raises the maximum length of a disc from one hour to two, sacrificing video quality if you opt to squeeze in more data.

This version also adds 20 themes for creating handsome navigation menus and a DVD map function that provides an overview of a project, which is useful for organizing a disc's presentation. DVD menus don't have to be static, as many of the included themes have optional motion menus. For example, with tropical fish swimming by or a flashing movie marquee, iDVD can produce a darn nice-looking disc.

A DVD that exploits its full potential means more than just video, so iDVD lets you incorporate audio from iTunes and GarageBand into menus. You can also include DVD-ROM data. Other nice touches include the ability to create slideshow discs and looped movies or slideshows.

Although you can edit a movie project on most Macs, the program can't burn discs unless the Mac has an internal SuperDrive. (Third-party FireWire drives won't work.) The program requires a 733MHz G4 or G5 processor, but it refused to install on my 400MHz G4 with a 1.25GHz upgrade card. I tested it on a 1.25GHz PowerBook G4 instead.

iTunes

iTunes didn't receive an update for iLife '04; its last major upgrade was more than a year ago. Maybe Apple's philosophy (for now) is not to mess with success. iTunes is a fabulous digital jukebox program, unparalleled by anything else on the Mac and (arguably) Windows. With it, you can listen to and organize your music collection, tune in to streaming Internet audio, rip CDs, and burn your own compilations. If you have multiple Macs, you can play songs from networked machines without fuss.

The program is also the gateway to the popular iTunes Music store, which now offers 500,000 downloadable songs from five major record labels, plus hundreds of independent labels. It's not just music, either; audio books and public radio programs have been added to the selection.

Putting It All Together

iLife '04 will definitely not run on your Mac Plus. The suite requires Mac OS 10.2.6 or later, 256MB of memory, and 4.3GBs of space. Most iLife components work with a G3 or later processor. GarageBand requires a G4 for full functionality. You'll need a DVD drive to install GarageBand and iDVD.

iDVD, iMovie, and GarageBand, the most complex programs in the suite, include helpful tutorials to get you started on the right foot. Each application in iLife stands strong on its own, but their integration is the icing on the cake. Working in iMovie, you can incorporate music from iTunes and still shots from iPhoto. You can easily export the songs you create in GarageBand for use in the other applications. ▲

by Kevin Savetz

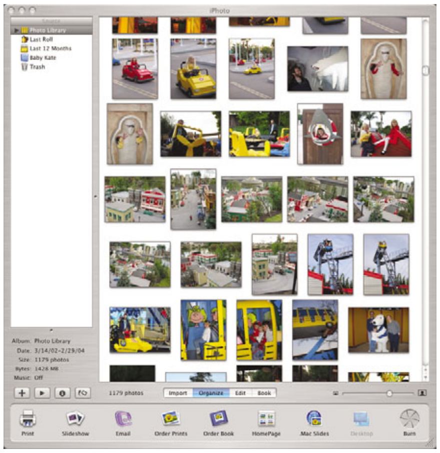

Album: Photo Library
Date: 3/14/02-2/29/04
Size: 1179 photos
Bytes: 1428 MB
Music: Off

iPhoto's ability to store, catalog, and manipulate your photos is marvelous.

> **A DVD that exploits its full potential means more than just video, so iDVD lets you incorporate audio from iTunes and GarageBand into menus.**

MAXIMUM PERFORMANCE

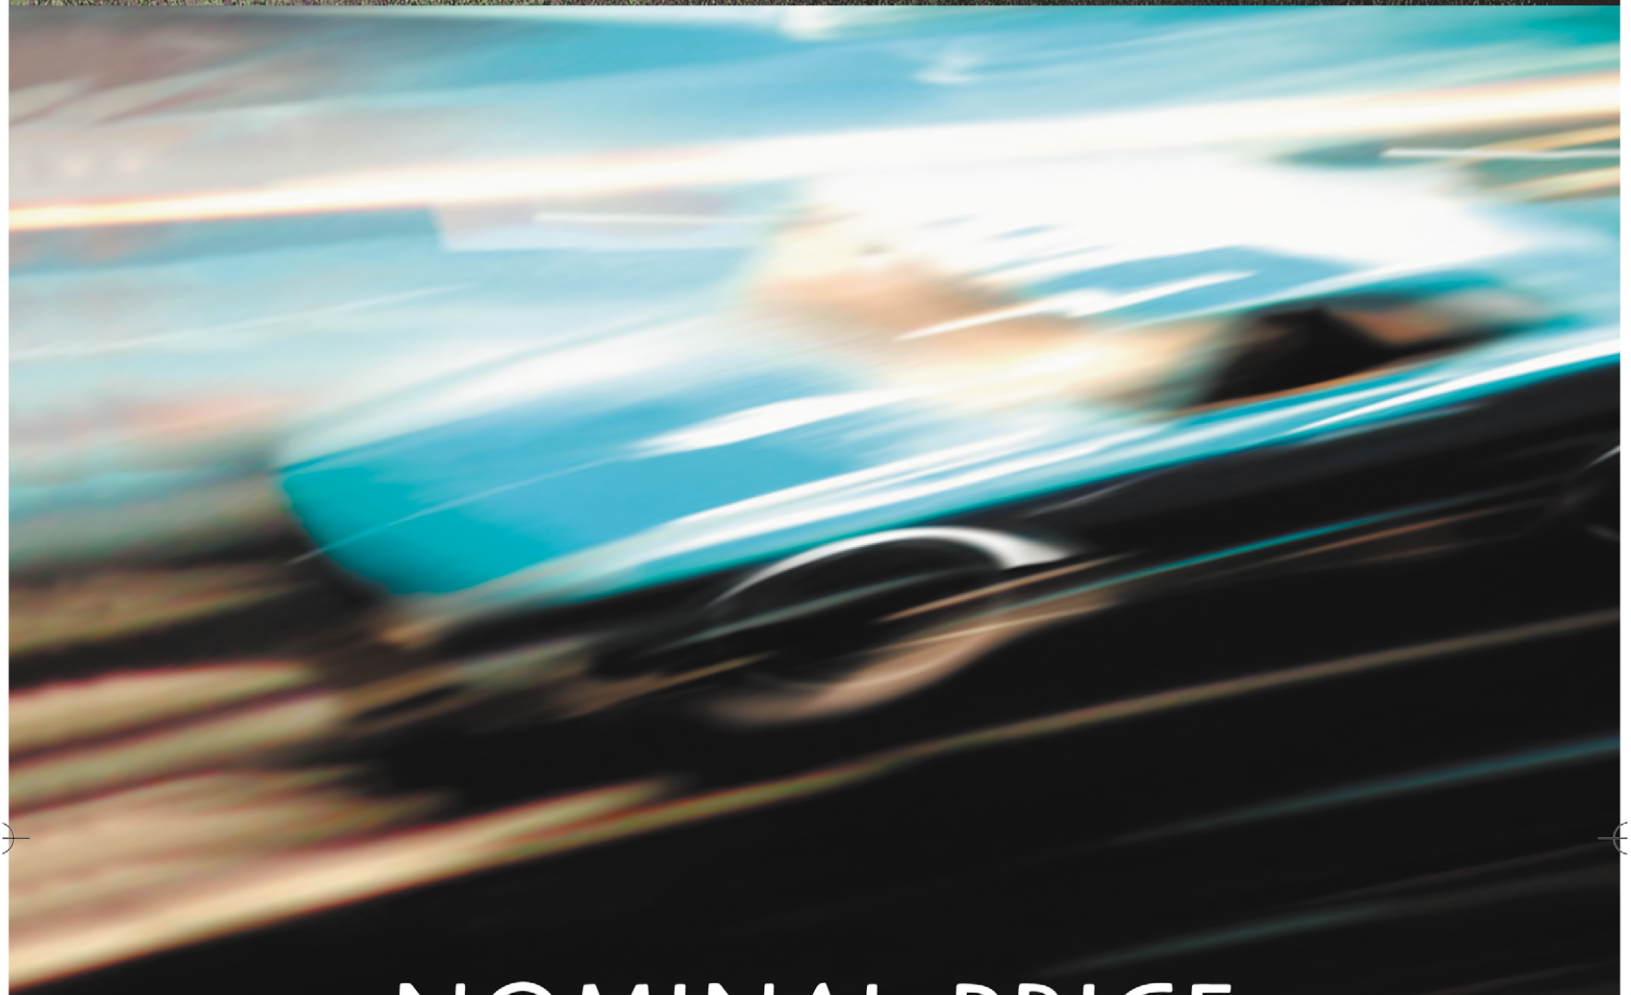

NOMINAL PRICE

Experience the new 8X DVD+/- RW Burner from AOpen

Now take advantage of the superior performance of AOpen's® new 8X DVD +/- RW drive at an amazing price!

AOpen – the world's leading PC component manufacturer – has pulled-out all the stops with this new feature-loaded, dual-format DVD burner. It not only writes both DVD+R and DVR-R media at true 8X speeds, it also serves as a handy, fast and convenient CD media writer and player. It may be a record-breaker, but there's one thing the DWR8800 DVD Burner is NOT, and that's a budget-breaker.

With extra goodies like Ahead's *Nero Burning ROM, InCD, VideoStudio 7,* CyberLink's *Power DVD* and 321 Studios' *DVD X Copy* (trial version) – you'll make movies, author professional-caliber DVDs and efficiently manage all your media storage needs – all with one nifty package.

It's fast. It's reliable. It's affordable. It's the AOpen DWR8800 8X DVD Burner. Maximum performance. Nominal price.

AOpen®

www.aopen.com
1-888-97-AOpen

© 2004 AOpen America Inc. All rights reserved. AOpen is a trademark and the AOpen logo is a registered trademark of AOpen Inc.
All other trademarks are the property of their respective owners.

Dialogue Box

by Chris Pirillo

Hijacked By Spyware

Chris Pirillo recently launched Lockergnome.net, a blogging network with features that geeks have been waiting for. This is an accompaniment to his wide-reaching Lockergnome.com, which delivers technology information to curious users over the Web, email, and the much-ballyhooed RSS. The fourth annual Gnomedex has been announced and will take place Sept. 30 through Oct. 2 in Lake Tahoe, Nev. No matter who you are or what you do, it seems that Chris is working on something to help you through the day. Did we mention that he's still doing that Rentmychest.com thing? Nevermind.

The worst of the Web continues to shine through murky browser windows, blocking your view of what's useful (and even fun) online. From pop-ups to pr0n, from searches to sales, the ilk that's creating spyware have latched on to all of us at some point. Shaking them off your leg is getting tougher by the day. The harder you kick, the tighter their grip becomes.

The vilest of the vile are the so-called search sites that take hijacking to an entirely new altitude of frustration. Without realizing what happened when you ran off to the loo, you were suddenly incarcerated. You can no longer visit the sites you prefer. You're no longer in charge of your own keyboard. You're a prisoner of your own computer. Trying to hit Google takes you everywhere but Google.com. Changing your home page manually and locking it down with a tool like Spybot Search & Destroy (safer-networking.org) is a seemingly fruitless effort. A desperate scan for viruses issues a clean bill of health, leaving you even more feverish. You wouldn't be the first to just reach for the power button and take a nap. Problem is, that only extends your sentence.

> **Remember: You're playing a virtual game of leapfrog in a minefield.**

Among the many resources that will hose the sludge from your prison cell windows is SpywareInfo.com. Mike Healan has been leading the battle against these evil forces, often branching into privacy and security issues that piggyback the payloads. Particularly educational are the forums, where victims expose the latest nasties, and more experienced exterminators offer assistance. Commonly discussed tools include HijackThis, which checks actual areas of attack rather than specific targets, and CWShredder, helpful in stomping the horrible search engine weeds taking over the Windows Registry.

Escaping is easier when you can see the wall, but evading repeated capture is an equal test of your patience. How many times have we seen spammers trying to sell antispam software? Does anyone really fall for that? The fact that spammers keep doing it confirms it's (sadly) profitable. Alanis Morissette would love the irony in this: Two of the most prevalent infestation methods come through commercial pop-up killing apps and spyware removal tools that inflict much more harm than help. All you're doing there is redecorating your cell, not gaining your release.

Spyware control is an area where the "little guy" is often more effective than corporate software creators. The scene changes very frequently and large organizations can rarely keep up with the latest threats, even though they may offer more features and somewhat refined interfaces. You might, then, consider using multiple products. Not every title is the best for every situation. Remember: You're playing a virtual game of leapfrog in a minefield.

Internet Explorer is typically the avenue taken by the cowards, and naturally, Microsoft has been shelled over these matters repeatedly. WinXP users will be happy to hear about upcoming changes on the security front. Service Pack 2, as of this writing, is in beta stages and set for release around midyear. With it comes the long-demanded pop-up blocking capability in IE. The built-in firewall also gets a revamp, providing protection against outbound traffic, as well as traditional inbound attack vectors. To download an advanced copy of the huge honkin' set of patches, feed "sp2predl" into your favorite search engine. Of course, you could always switch to a Mozilla flavor, Opera, or (heaven forbid) Lynx.

We'll undoubtedly hear cries of foul behavior toward third-party firewall vendors (ZoneLabs' ZoneAlarm, Agnitum's Outpost, Tiny Software's Personal Firewall), but tough cookies. There's plenty of room for additional features, and the way that Microsoft is crafting the changes in SP2 is the single most respectful act toward users and system administrators I've seen to date. While pop-up killing in IE is enabled automatically (a welcome default), smart choices are offered for automatic software updates and the revamped Windows firewall. No default options are predetermined for you at each initial configuration dialog. Not only do users have little excuse for how their PCs are configured, they'll likely pay more attention to what's being asked of them and soak up a bit of education along the way.

The battle for your PC rages on with no end in sight. Multiple layers of armor and a familiarity with removal tools before the impending assault will keep you from screaming yourself hoarse. ■

You can dialogue with Chris at chris@cpumag.com.

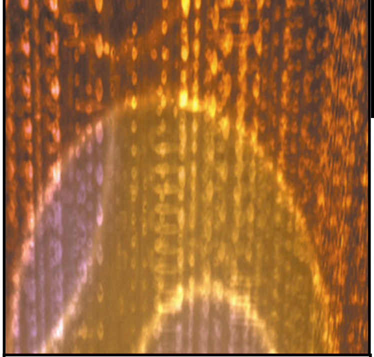

Open Sauce

by Pete Loshin

Freaky Wiki

Pete Loshin, former technical editor of software reviews for Byte Magazine (print version), consults and writes about computing and the Internet. He also runs http://www.linuxcookbook.com. He owns shares of both Microsoft and Red Hat and believes that Windows isn't for everyone, but neither is Linux.

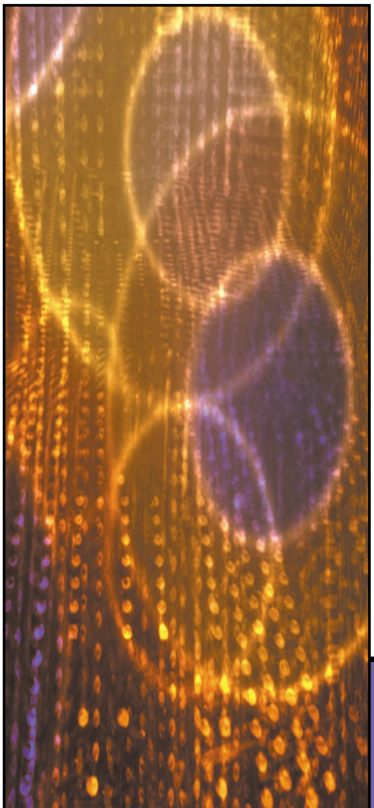

Remember Gopher? Unless you were wired before 1994, probably not. Gopher (see RFC 1463, "The Internet Gopher Protocol") was a text- and menu-only hypertext system. Big in the early 1990s, it went belly-up with the Web explosion.

Wiki ("wee-kee" or "wicky"), an application that runs on the Web, reminds me of Gopher—but in a good way. Ward Cunningham, the father of wiki, wanted "The simplest online database that could possibly work." First implemented at the Portland Pattern Repository (c2.com/ppr) in 1995, wiki, according to the "What is Wiki" page at wiki.org, is described as the following:

- ". . . a piece of server software that allows users to freely create and edit Web page content using any Web browser. Wiki supports hyperlinks and has a simple text syntax for creating new pages and crosslinks between internal pages on the fly."
- "Wiki is unusual among group communication mechanisms in that it allows the organization of contributions to be edited in addition to the content itself."
- "Like many simple concepts, 'open editing' has some profound and subtle effects on Wiki usage. Allowing everyday users to create and edit any page in a Web site is exciting in that it encourages democratic use of the Web and promotes content composition by nontechnical users."

> ". . . wiki will be going strong in 10 years . . ."

Good wikis are owned by everyone and no one. Anyone can contribute. Helpful and useful contributions are verified by others. Unhelpful ones are edited out. Wiki reminds me of Gopher for various reasons, including:

- Many wikis are about wiki technology, history, and philosophy, just like most Gopher sites and early Web sites, for that matter.
- Wikis often go in circles, with many pages pointing at each other and/or repeating information.
- Wiki pages often have so many links they look like menus, and few seem to have many graphics or other non-text content.

Unlike Gopher, some form of wiki will be going strong in 10 years, for numerous reasons, including:

- Wiki invites involvement and improvement. See an error and fix it. If there's something missing, go ahead and add it.
- Wikipedia (http://en.wikipedia.org), an extraordinarily useful wiki, is a free online encyclopedia available in dozens of languages under the GNU Free Documentation License. The English version has hundreds of thousands of articles, and it rivals traditional references in terms of article length, number of articles, and immediacy.
- Wiki is the best approximation to date of Vannevar Bush's original Memex. See www.theatlantic.com/unbound/flashbks/computer/bushf.htm for the 1945 *Atlantic Monthly* article in which Bush describes his hypertext vision for the future. Or compare Wikipedia to Douglas Adams' "Hitchhiker's Guide to the Galaxy" from the book of the same name.

As a fast-and-dirty Web collaboration tool, wiki may eventually dominate the long-crowded corporate groupware category. The only client needed is a Web browser; wikis can be set up on any Web server. As long as know-it-alls with Web access have too much time on their hands, general-knowledge resources such as Wikipedia will only get better.

Updates & More

In February 2003, I mentioned the open-source hardware Simputer (simputer.org) project. Amida, the first commercial Simputer, is available now from PicoPeta Simputers (amidasimputer.com) starting at $240. At press time, it isn't shipping outside of India, but it plans to soon.

When haven't I ranted about not tossing old PCs simply because they don't run the latest Windows version? If you have an obsolete PC, see RULE (Run Up2date Linux Everywhere) at rule-project.org/en for help installing a modern Linux distro.

A year ago, I wrote about the GIMP as a free alternative to PhotoShop. The long-awaited version 2.0 is finally out (gimp.org). Check it out.

I'm a bit too old to have gotten hooked on computer gaming, but I'm going to try this year. The first thing I'm going to try is the Cube (cubeengine.com), a multiplayer/single-player FPS. It's optimized for Linux, can run on modest hardware, and you can even develop your own games with it. ■

Get saucy with Pete at pete@cpumag.com.

Can CAN-SPAM Really Can Spam?

New Laws Are No Match For Unwanted Email

You would think that email spammers would be closing their shutters and running to their hideouts by now. After years of users moaning about the torrents of unwanted email cluttering their inboxes, legislators and prosecutors finally took action this year, slamming the spammers with a flurry of new laws and court cases.

In January, the federal CAN-SPAM Act went into effect, which imposes criminal penalties on egregious inbox violators. In the spring, AOL, Microsoft, EarthLink, and Yahoo! joined forces to sue hundreds of suspected spammers in several states. AOL made an example of one convicted spammer by raffling off his $47,000 Porsche, which was seized in the arrest. Even home repair guru Bob Villa was caught in the spam dragnet when a California ISP sued his home how-to site (www.bobvilla.com) and its email marketing company for violating the new federal laws.

The CAN-SPAM (Controlling the Assault of Non-Solicited Pornography and Marketing) Act of 2003 was signed into law in late 2003 and went into effect Jan. 1, 2004. It was the result of six years of haggling, lobbying, and aborted legislative attempts to impose a single federal law concerning abusive email practices. Before this time, the Federal Trade Commission could only prosecute spammers by using old laws pertaining generally to deceptive advertising. CAN-SPAM was the first time a federal law went into detail about what deceptions and violations specific to email were illegal and criminally prosecutable. That means spammers can now see jail time. The law requires that unsolicited commercial email messages be labeled with the accurate sender, that they not use deceptive subject headers, that they include unsubscribe options, and that the sender's physical address is included. In addition, the law gives the Federal Trade Commission the power to create a "do not email registry," much like the famous Do Not Call Registry, which might let users sign up once to opt out of all email solicitations.

But if you are like most of us, you probably haven't noticed any downturn in the number of missives in your inbox. In a survey of Internet users earlier this year, the Pew Internet and American Life Project found more than half of users detected no reduction in the amount of spam they received since CAN-SPAM went live. According to email security provider Brightmail (www.brightmail.com), the percentage of email its systems were filtering out for clients as spam rose from 60% in January 2004 to 63% in March 2004. Jupiter Research says that the average online user received 3,920 unwanted emails in 2003, and by 2008 they will get 6,395. The bottom line is that almost half a year after the law went into effect, and despite high-profile lawsuits, spam actually seems to be increasing, and very few industry analysts see the new push working.

Spam Un-Canned

"I don't think [new laws] are going to have any effect on malicious spammers," says Sara Radicati, president and CEO of tech consultant company The Radicati Group (www.radicati.com), which recently studied the email problem. She makes the dire prediction that up to 142 billion spam emails a day will be sent in 2008, constituting 74% of all messages. A little more than half of emails today are considered spam. "Spam has rapidly transformed from a mere nuisance factor to a worldwide epidemic that is crippling organizations and posing serious security threats to email and network infrastructures," she says.

Worldwide Spam Growth Forecast For 2004 to 2008 (in billions)

A recent study by The Radicati Group revealed a rather dire outlook for the spam problem in the future.

	2003	2004	2005	2006	2007	2008
Total Worldwide Email Messages/Day	32	68	86	109	143	191
Total Spam Messages	15	35	53	73	102	142
Percent Spam	45%	52%	62%	67%	71%	74%

Source: Radicati Group

This is costly. At companies especially, spam clogs email servers, slows systems, and ultimately requires the enterprise to expand capacity or buy costly filtering systems. Radicati says that companies and individuals already spend nearly $1 billion a year fighting spam, which could increase to $1.7 billion by 2008.

Even legitimate marketers decry the cost of spam because they, too, need to invest in costlier ways of ensuring their permission-based emails comply with new laws and even get through the more sophisticated filtering mechanisms in place. In fact, when the CAN-SPAM law took effect in January, its main consequence seems to have been frightening off mainstream advertisers. "Most of us stayed dark in January," says email marketer and analyst Larry Chase, publisher, Web Digest for Marketers (www.wdfm.com). "I think [CAN-SPAM] dampens legitimate marketers. They wanted to see what was going to happen. We're into a new frontier."

Chase says that legitimate marketers have a reputation to protect, and they really don't want customers to be upset with the brands they advertise. Not only do they need to ensure that they stay inside the new laws, but they need to make sure customers and ISPs perceive them to be working inside the laws. Chase admits that CAN-SPAM is a "good start" that empowers prosecutors to go after the worst offenders, but its effect on the level of spam in our inboxes won't be felt for a long time.

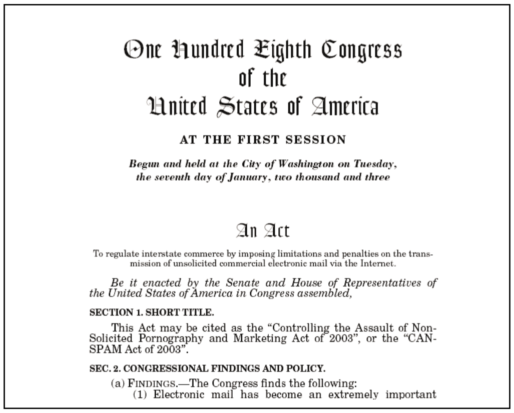

One Hundred Eighth Congress
of the
United States of America

AT THE FIRST SESSION

Begun and held at the City of Washington on Tuesday,
the seventh day of January, two thousand and three

An Act

To regulate interstate commerce by imposing limitations and penalties on the transmission of unsolicited commercial electronic mail via the Internet.

Be it enacted by the Senate and House of Representatives of the United States of America in Congress assembled,

SECTION 1. SHORT TITLE.

This Act may be cited as the "Controlling the Assault of Non-Solicited Pornography and Marketing Act of 2003", or the "CAN-SPAM Act of 2003".

SEC. 2. CONGRESSIONAL FINDINGS AND POLICY.

(a) FINDINGS.—The Congress finds the following:

(1) Electronic mail has become an extremely important

The CAN-SPAM Act of 2003 is the first federal legislation specifically written to regulate email marketing and apply criminal penalties to violators.

Why CAN-SPAM Can't

Many prosecutors and Internet analysts complain that the CAN-SPAM law simply does not have the teeth even to nibble at the spam problem. The law doesn't prohibit unsolicited emails; it simply gives these messages a set of rules to obey when spamming millions of inboxes. Moreover, many state legislators and prosecutors are angry that the CAN-SPAM Act pre-empts state proposals that carried much stricter rules and penalties, and it does not allow for individuals to sue email abusers. CAN-SPAM forces email marketers only to give users a way to opt out from getting more messages.

Several states, including California, have legislated an opt-in approach that bans unsolicited emails altogether. Companies are not allowed to email someone unless the consumer has an existing relationship with the company and tacit permission to contact the user. Calling CAN-SPAM an ineffective half-measure, California State Senator Debra Bowen said at its passing, "The bill doesn't can spam, it gives it the congressional seal of approval. Spam is spam. There's no such thing as 'good spam' and 'bad spam.' "

CAN-SPAM is hopelessly vague, critics argue, because it never defines acceptable email labeling nor does it specify that unsubscribe options must be prominent and easy to activate. Some entire categories of senders, including political, nonprofit, and religious groups, are exempted altogether. As a result, some of the very people who would be prosecuting spammers, the National Association of Attorneys General, urged lawmakers not to enact it. "The bill creates so many loopholes, exceptions, and high standards of proof that it provides minimal consumer protections and creates too many burdens for effective enforcement," the group said in a statement to Congress.

And after all of these years of discussion, there is still an ongoing disagreement among marketers, users, and corporate IT managers about what constitutes spam. "There are many, many definitions," says Radicati. "We still feel it is a gray area." Professional email marketers like to define spam as "unsolicited" or non-permission-based email. The problem is that "permission-based" messaging is very widely defined. Users often feel that "permission" involves specific approval for one commercial entity and one alone to contact them. For many marketers, "permission" often encompasses massive mailing lists that are traded and sold among them because at some point those consumers gave someone "permission" to email them and then failed to regard the fine print disclosure that the original company could pass their names on to others.

Moreover, enterprise IT managers often have their own idea of spam as just about any consumer-oriented commercial message that comes to a worker's inbox. "As an individual I may be happy to get a mailing from an affinity group I belong to, but a company thinks it is a waste of my time," says Radicati. "There is a big difference between what a company considers spam and what an individual considers spam."

Ultimately, spam is not going to go away because it is simply too easy and lucrative for spammers to operate and cover their tracks. When Canadian email security firm Vircom (www.vircom.com) interviewed three professional spammers early this year, it discovered that most could be up and running for $1,500. Many of them are individuals who work under contract to marketing firms or adult sites and can make between $150 to $2,000 per spam campaign, up to $6,500 a week. Apparently all it takes is a single computer; an ISP; software that harvests email addresses from forums and newsgroups; and spamware, which is designed to send bulk emails in ways that avoid detection from the host ISP and filters.

One 29-year-old spammer, Matt, told Vircom interviewers that he considers it a legitimate business. "We are marketing and promoting items that large companies, Fortune 500s in many instances, pay us to market. It's a typical big business; they go after the little guy but let the big businesses off the hook." Vircom also found that the spammers it interviewed were not concerned about CAN-SPAM because the wording of the law is vague and it may not be easily enforceable.

The Solution Problem

Antispam technologies on the receiving end, either deployed by ISPs or individuals on the desktop, are having some effect on spam levels. In a survey of corporations, Radicati found that while many companies suffered 52% spam rates before installing spam filters and other counter-measures, spam percentages dropped to 17% with filters. For consumers, however, the spam share dropped only to 30%. According to Arbitron/Edison Media Research, 50% of consumers now employ some sort of spam-blocking tool, up dramatically from just 34% in 2003. Nevertheless, many complain

The Big Spam Killers?

In the end, many say that fighting spam requires a high-level tech solution deployed across the Internet. To that end, the Internet Research Task Force, a group of respected engineers who research and recommend new standards for the Internet, agreed to take on the spam issue this year by forming the ASRG (Anti-Spam Research Group; http://asrg.spam). Over the coming months, the group will review the problem and evaluate a host of different solutions, including some of those outlined below.

Caller ID. Proposed by Microsoft, this system requires legitimate marketers to submit the IP address of their outbound email servers to the DNS (Domain Name System). On the receiving end, email clients could then compare the actual route an email took to the inbox against the DNS listing for that route to determine if the message is indeed coming from the domain claimed by the sender.

Domain Keys. Implemented by Yahoo! and proposed as an Internet-wide solution, this approach puts an encrypted key for a legitimate sender in the header information of an email, which can only be unlocked by a corresponding code registered with the DNS. The recipient's email client gets the assigned key for the purported sender from DNS to unlock the message. If the DNS code does not match the key in the message, then the email probably came from a spoofed source.

Sender Permitted From. The SPF (also know as Sender Policy Framework) approach is backed by AOL and resembles other authentication in that it has ISPs register their email servers with DNS so that other ISPs can check the claimed sender of an email against its actual IP source and determine whether the sender is being faked. Unlike Caller ID, however, this method checks an email's source from its "wrapper," the information exchanged between servers before a message is actually sent. ▲

Infinite Loop

Emanic Monday

If you start your week online, you're not alone. OneStat.com recently determined that Monday and Tuesday garner 30.14% of the week's online activity worldwide. (Monday takes the lead with 15.43% to Tuesday's 14.71%.) Users take a break Sunday, accounting for only 12.52% of the week's activity. ▲

The Dirty Dozen Of Spam Sources

Internet security firm Sophos (www.sophos.com) monitored and tracked to its sources spam sent over the Web during a two-day period in February 2004 to discover the top 12 spam producers in the world.

1. United States	56.74%
2. Canada	6.80%
3. China (& Hong Kong)	6.24%
4. South Korea	5.77%
5. Netherlands	2.13%
6. Brazil	2.00%
7. Germany	1.83%
8. France	1.50%
9. United Kingdom	1.31%
10. Australia	1.21%
11. Mexico	1.19%
12. Spain	1.05%
Others	12.23%

Source: Sophos

that current solutions cause as many problems as they solve. Almost half of corporate Radicati's respondents said they often didn't have access to quarantined emails, and many others complained about the absence of customization on their corporate emails.

While many marketers claim that anti-spam legislation has little effect because so many email violators come from overseas, the statistics suggest otherwise. The United States is responsible for more than half of the world's spam, with most other sources far behind. (See "The Dirty Dozen Of Spam Sources" chart.) If technology can gain a foothold in this country, it might go a long way in eliminating the worldwide problem.

Many feel that the real solution to spam is years away and that it will take the form of a broad technology innovation that becomes standard. The major players such as Microsoft, Yahoo!, and AOL have each proposed new techniques for sending and receiving email that help authenticate a sender as legitimate or trusted and block off the usual routes spammers use to deceive ISPs and users. Microsoft has proposed an email version of caller ID, and AOL is behind a technique called SPF (Sender Permitted From). Legislators can pass as many laws as they like, but many point out that criminal penalties and prosecutions have not dissuaded virus and worm writers, so they won't impress spammers. "I think the big quantum leap in dealing with all of this will be the set of new proposals like Caller ID and SPF that attempt to verify authorization," says Radicati.

Nevertheless, spammers are already charting the next frontier of intrusive and annoying messages by getting even more personal. SPIM, or unwanted messages via instant messaging, may triple in 2004 alone, from 400 million a day last year to 1.2 billion this year, says Radicati. On yet another front, cell phone spam started cropping up last year as some daring marketers started using mobile text messaging as a way to bulk mail consumers. And as many cell phone users already know, some of us pay for every message we receive, as well as send. CPU

by Steve Smith

More than a monitor,
it's a window to your memories.
BenQ

View digital photos instantly — no PC required. The BenQ FP591 and FP791, featuring Digital Photo Frame (DPF), allow you to view digital photos instantly just by inserting your camera's memory card. No cables, downloads or hassles. Just pure enjoyment. And the image quality is every bit as spectacular. A+ panel quality (ranking at the top of the industry's stringent standards), lightning fast 16ms response time, ultra-high brightness/contrast ratio, and SRS® stereo surround sound combine for a display you'll cherish as much as your fondest memories.

PC MAGAZINE 2002 BEST OF COMDEX FINALIST
Nov. 20, 2002 BenQ FP591
Hardware, Peripherals

For more information, visit www.BenQ.com/2003/SmartComputing-FP791
To contact us, visit www.BenQ.com/ProdInfo or call (866)700-2367

A+ PANEL QUALITY

©2003 BenQ America Corp. All rights reserved. PC Magazine "Best of Comdex" FINALIST Award Logo is a trademark of Ziff Davis Publishing Holdings Inc. Used under license. All trademarks and registered trademarks are the property of their respective holders.

BenQ
Enjoyment Matters

Coder's Corner: XML

Schema Languages For XML

Over the past few months we've introduced DTDs (document type declarations), which are used to (among other things) define specific types of XML documents. Using a DTD, a validating XML processor can check that an existing document is consistent with the rules defined in the DTD and reject those that aren't. Similarly, an XML editor can use the rules in a DTD to enforce consistency as a document is written. With a validating editor, an author simply can't put text or elements in the wrong place; the editor won't let him.

However, as we noted last month, DTDs have several weaknesses when it comes to XML. One weakness is that DTDs have no innate understanding of XML namespaces. Another is that DTDs can place only simple constraints on element content or attribute values. For example, DTDs can say that an attribute value must be a string (CDATA), an ID (ID, IDREF, or IDREFS), a name token (NMTOKEN or NMTOKENS), or a list of allowed tokens, such as value1 | value 2 | value 3. But DTDs can't say that the value must be a positive integer, a date (subject to a particular rule for writing dates), and so on. The situation is similar for element content.

In general, a schema provides a set of rules that models a structure and defines constraints on that structure.

Unfortunately, it's these sorts of constraints, or type validation rules, that are critical when using XML to encode data in a general-purpose programming environment, such as for a database table, a data structure, or a Web service.

The experts who developed XML 1.0 recognized these weaknesses early on and gave early consideration to simply omitting DTDs from the formal specification. However, this idea was quickly discarded for two simple reasons. First, the designers knew that including DTDs would give XML a big head start, as there was a lot of existing DTD-based technology and data (inherited from SGML) that could be quickly migrated into XML. Second, there was no answer to the question: If not DTDs, what else? Indeed, there were strong arguments that it was too soon in XML's life to know what a replacement should look like and what new features it should have. So DTDs were included, with the expectation they would soon be augmented or replaced by a technologically stronger successor.

To say this process took longer than expected is an understatement; it was over three years before the W3C (World Wide Web Consortium) defined a new official XML schema language. We'll address what happened and why. But first, we'll look at what schema languages are really for.

What Is An XML Schema Language?

In general, a schema provides a set of rules that models a structure and defines constraints on that structure. For example, a relational database schema defines the database's structure, including the types and sizes of tables, the fields supported in each type, any relationships/constraints between different fields, and more. Anyone who has built an Access database has designed a schema. The schema is created by the graphical design of the database that a user builds before inserting any data. A schema language is just a way of writing the schema instructions so software or users can unambiguously understand them.

For XML, a schema language must be able to model the structure and constraints for types of XML documents. This is equivalent to the goal for database schemas, but the scenarios by which the schemas will be used add a number of novel requirements. The most important is universality;

the language must be universal so everyone can understand and process the schema with precisely the same results. (Database schema specifications are often restricted to a specific database type or vendor.) This is part of the core XML philosophy: XML data should be universal and processed/interpreted in the same way by any consumer of the data. But other requirements—mostly focused on making sure the language was useful to those working with XML (databases, documents, data modelers, etc.)—placed a heavy burden on designers.

In early 1999, these and other needs were summarized in a W3C requirements note. Some core goals were that the new XML schema language be:

- Better than XML DTDs (a superset of functionality)
- Written in XML
- Self-describing (there is a schema describing the schema language)
- Usable on the Internet
- Integrated/coordinated with other core XML specs

Right out of the gate, there were several prototype XML schema languages that presented alternative approaches for accomplishing these goals and satisfying the requirements.

Why Did It Take So Long?

With all these examples available, why did it take so long to complete a specification? The answer comes in three different parts.

The first problem lay in the lack of foundational XML technologies needed to support an XML schema language. For example, an XML schema language needs a way of referencing potential parts of a document so that it can define constraints on those parts. This is the role of XPath (and XPointer), but XPath was only standardized in late 1999.

The second perceived problem was the lack of a clear and unambiguous model for the data structure underlying an XML document. Such models had already been implicitly defined in other XML specs that interacted with or referenced XML data structures, such as the XML DOM (document object model) and XPath. However, the models underlying DOM and XPath are different. For example, the DOM sees CDATA sections as distinct parts of an XML document, but XPath doesn't. The W3C schema working group felt it was important to have up-front consensus on a universal abstract data model (called the XML Information Set, or Infoset for short) so it could build a schema language consistent with it.

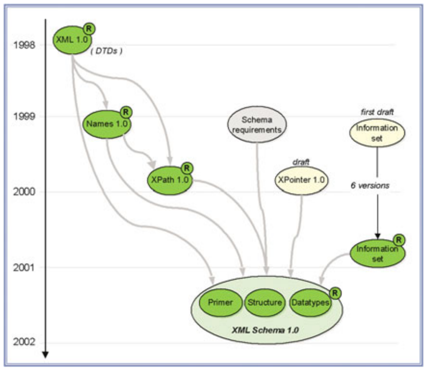

This illustration shows how the XML schema specifications depend on other official XML standards. You also see rough timelines for the creation of these documents. The R inside a circle means the associated item is a W3C recommendation, or a formal specification. The gray circles are W3C notes. The yellow circles represent working drafts of a specification.

Last was the sheer scope and complexity of the task. This lay in two areas: the people involved (from academic experts to large corporations) and the breadth of issues in scope (from traditional document modeling to semantic Web design to database models to Web service design). This led to lots of wrangling over the best approach to take, plus debate over the scope of issues that should be included. Although the requirements we've noted constrained the scope somewhat, the task was still daunting and large.

The XML Schema Specification(s)

The W3C XML schema specification was finalized in May 2001. Because of the complexity of the work, the specification was broken down into three documents.

Datatypes. This is a specification for a type system for XML schema that spells out how schema authors can extend the system. In essence, the type system applies to string values, attribute values, or elements that have only text content.

Structure. This specification is for defining structural rules for XML documents, or structures with elements inside elements or with attributes attached to elements or both. This includes support for namespaces and various mechanisms for extending or inheriting already defined rules (based partly on ideas from object-oriented design).

Primer. This explains how the formal specifications in the structure and datatypes documents are used to create schema definitions using informal language and examples. (It's probably not a good sign that the specification needs a primer to explain how it works.)

The datatypes and structure specifications are each over 100 pages long and are almost impossible to understand without a good understanding of the XML information set. The primer is better and is the place to start when learning about XML schema. Although XML schema is complex, much effort has been taken to implement it in XML tools from XML editors to Web service infrastructure. Today the language is quite widely supported. Next month we'll look at some simple XML schema examples to give a flavor for how it works.

XML schema isn't without competitors, however. Early on, several notable XML experts believed that XML schema was heading down a too-complex and perhaps incorrect path. These experts, including James Clark, author of the XSLT specification and co-author of XPath, decided to try other approaches. With fascinating names like TREX, Relax, and Schematron, these approaches led to some different and, often times, easier-to-understand schema languages for XML. We'll look at these alternatives in future articles.

by Ian Graham

(For a list of useful books and other online references on XML schema, go to www.utoronto.ca/ian/articles/jun04.)

The Department Of Stuff

by Rob "Cmdr Taco" Malda

uribl.txt

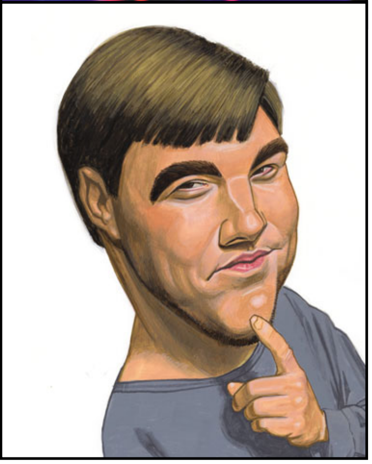

Rob "CmdrTaco" Malda is the creator and director of the popular News for Nerds Web site Slashdot.org. He spends his time fiddling with electronic gizmos, wandering the 'Net, watching anime, and trying to think of clever lies to put in his bio so that he seems cooler than he actually is.

As the volume of spam continues to grow at rates that baffle the imagination, the escalating war between the spammers and the tortured masses continues to rage. Over the past few years, two major developments have changed the battle completely: blacklists and Bayesian filtering, and a third is about to arise.

By now everyone is familiar with Bayesian filtering—from SpamAssassin to the default Mail.app included with OS X, anyone who spends any serious time reading email uses a Bayesian filter to sift out all the obvious spam. Unfortunately, these filters worked too well. Spammers have now begun including intentional misspellings, random words, or just plain made-up gobbledygook in their pitches in an attempt to squeeze past your filters.

I can't say this surprises me. Spammers are a desperate lot. They make their profits only because they can send their pitch for free. . . . Of course, the cost to ISPs is immense, to say nothing of the cost to you and me, both in the minutes of each and every day wasted by deleting spam and in the increased costs our ISPs must pass along to us while they deal with billions of extra emails each day.

Every piece of spam you get has one thing in common: a URL.

Unfortunately telling the difference between a spammer and someone with a tendency to mipsell wrods is a difficult task at best. Which is why people are now looking back to the old blacklist as the next step in stopping spam. But where a blacklist in 2001 meant blocking mail servers, in 2004 we're not analyzing the domain names sending the spam . . . we're looking at the domain names of the domains you are being asked to click. Blacklisting domains from sending email is a hugely error-prone mechanism. Many spammers use large ISPs to send their mail, and blacklisting those ISPs catches thousands of innocent email users in the crossfire.

But what does every piece of spam have? No, not Viagra, nudity, or offers of lower mortgage rates. Every piece of spam you get has one thing in common: a URL.

Those http:// things are the purpose for every single spam you get. There are a thousand ways to try to sell you something, but Mr. Spammer isn't going to get any money unless you click that innocent blue chunk of underlined text. 100% of my spam contains a URL, but only 50% of my non-spam does.

Historically, blacklists have used several different technologies for maintaining their list of banned domains, but the simplest method is to make use of one of the oldest and most reliable of protocols: DNS. Instead of querying a name server to resolve IPs and Names, URI Blacklists resolve known spam domains. Because DNS technology is simple, scalable, and fast, implementation is simple: parse the URLs from your email and do a DNS lookup using a name server that knows not names and IPs, but rather only if the domain in question has been used to promote spam.

This new technique is not without its problems. Like the junk words that spammers began inserting into their messages to combat Bayesian filtering, malicious marketers will begin using redirectors to disguise the final destination of the URLs you click, perhaps even using infinite loops of redirection in order to confuse would-be spam checkers. Blacklists will need to account for this by resolving these URLs to determine what the ultimate destination of each URL is. This is a nontrivial proposition.

Spammers will also include random innocent URLs in their messages to disguise the true nature of their message. It will be important for the maintainers of blacklists to figure out ways to tell these innocent sites apart from the true spam destinations.

Ultimately, if a spammer is going to get you his message, he needs to eventually deposit you on a page that has his complete pitch. And this pitch needs to be sent a million times in order to get 10 people to bite. It doesn't take very many users running URL blacklisting software for these servers to become very clear.

You can learn more about these ideas by reading www.surbl.org, the leading site in developing this new step in the war on spam. The site includes SpamAssassin 2.63, a 3.0 plug-in, and plug-ins for a few MTAs. ■

Please don't send any spam to malda@cpumag.com.

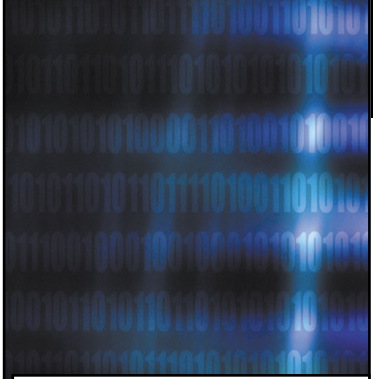

Forward Slash

by Joan Wood

Play Phone

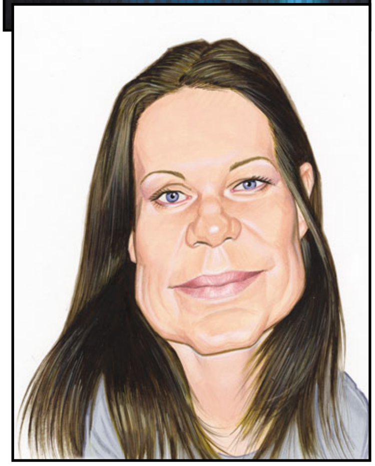

Starting as gopher for the Emmy-winning team that pioneered live in-car TV cameras for the Indy 500, Joan became an independent video/sound engineer, technical director, and producer. Playing with Reality Engines and motion platforms led to co-founding Xatrix Entertainment, where she produced the two Cyberia games. Before 3D acceleration was trendy, she formed Mango Grits to develop hardware-only game Barrage for Activision. Since cashing out from SharkyExtreme.com, where she was co-founder and managing editor, Joan has retired.

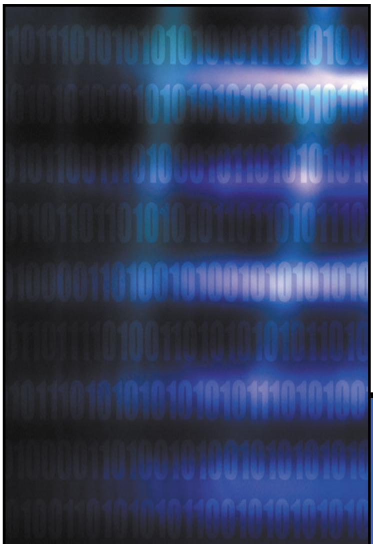

It was obvious when writing about game development for mobile devices on this page a year ago (Forward Slash page 85, *CPU* May 2003), that the state of the industry was still pretty green. Handset makers jostled over features, service providers wanted to be publishers, developers were trying to evaluate the level of commitment and savvy of each participant on a case-by-case basis, and everyone was looking for the business model that could spread the revenue throughout the food chain.

What a difference a year makes. No, we aren't playing Far Cry 3G, but things have . . . progressed. Some big players have taken seats at the table, and they do know a bit about creating platforms for games and assisting developers in building for them. On the software side, OpenGL and DirectX APIs are represented with special handheld versions, and on the hardware side, both ATI and NVIDIA showed up, and each has brought a very special gift: developer relations.

No, we aren't playing Far Cry 3G, but things have . . . progressed.

The graphics giants have begun selling some nice little chips to the handheld industry, as each has perceived enormous growth potential. Both started with the multimedia/video processing route, first introducing chips that do that part well. There has been some licensing of 3D technology to existing handheld chip veterans going on, but both are saving the big stuff for later. Even without 3D, things like dual-display capabilities have migrated from your PC to your two-screen flip phone, and the onboard still and video camera specs keep getting more impressive.

ATI sees that it is necessary to start being involved now, but not in a big splashy way. They will reach out to developers and publishers, provide education and presentations at events, participate in standards bodies, and generally work toward pushing the whole train along. ATI realizes that at some point, if they don't do anything to stimulate creation of content, there won't be any need for new graphics hardware. But until the whole thing goes mainstream, it will not be very interesting to big developers and publishers, and they will continue to farm out porting their PC and console titles to handheld. A minimum graphics platform is needed to start down that path, but once a high volume of chips with the same core goes into a lot of devices, watch out.

The existing game distribution model works because current 2D games fit in small enough files (about 250KB) that the carrier just charges customers to download them for a fee. The developer gets its cut, the carrier gets its cut, the customer gets a game, and everyone is happy. 3D game file sizes are not practical to download onto a phone, so distribution will have to be through some kind of media or the Internet, bypassing the carrier, which is a fundamental no-no. The carrier owns the customer relationship, so nobody wants to alienate them, and this is where the current business model is floundering a bit. There are proposals to get past this logjam, some that even impact game design. Instead of every game having crates or a plot where you save the world, every handheld game will have your character's first action be to place a phone call so the carrier can get his cut. That solution probably won't make the cut, but it would be a funny one, and it points to a need for some party to take a leadership role in hammering this stuff out.

As the newly hired Director of Handheld Content at NVIDIA, Brian Bruning (yes, the same guy who was head of developer relations at 3Dfx) is well-suited for the task and he spearheads NVIDIA's business strategy for taking NVIDIA 3D to handheld with great content, including talking with carriers, manufacturers, game developers, and publishers about profit distribution mechanisms to keep the carriers happy. NVIDIA's calculated move to "touch every pixel on the planet" has definitely spread to handheld. There are, apparently, more untouched pixels hiding out on mobile phones than have previously been discovered lurking near your PC, and NVIDIA has now committed the resources, including their developer relations, developer education, and TWIMTBP marketing program in an attempt to touch them, as well. ■

Dial joan@cpumag.com.

by Chad Denton

Road Warrior

The Future Of Wireless, Sendo vs. Microsoft, Unwiring At The Park & More From The Mobile Front

Motorola's A840 includes support for both GSM and CDMA networks, making it compatible in most of the world. Oh yeah, it also has a 1.2-megapixel digital camera.

Moblogging: A Use For Your Camera Phone

As long as you're spending hours on the road for work, why not share some of the sites with friends, family, or even an eclectic assortment of random fans? Moblogging, or mobile blogging, is becoming an increasingly popular way for camera phone users to do just that.

If you want to start your own moblog, a number of Web sites can help. One of the most popular is TextAmerica (www.textamerica.com). Once you register, you can send pictures directly to TextAmerica, and they will be available online in seconds. Other popular sites include SnapNPost (www.snapnpost.com) and Yafro (www.yafro.com), the brainchild of the Hot Or Not founders. Moblogs often include forums where individuals can share their thoughts and interact.

Mobloggers likely found a couple of interesting developments at this year's CTIA (Cellular Telecommunication Industry Association) tradeshow in March. Numerous handset manufacturers were showing off a new generation of megapixel camera phones. Motorola's V710 is a CDMA phone with Bluetooth. Even more intriguing is Motorola's A840, which is one of a few phones that supports GSM and CDMA technologies, plus it features a 1.2-megapixel camera. Sony Ericsson's S700 and Nokia's 7610 also include megapixel cameras.

The higher-resolution images can dramatically improve picture quality, but larger file sizes are going to need a faster data network. Thus, news on EVDO (Evolution Data Optimized) at CTIA probably turned a few heads in the moblog crowd. Verizon Wireless has already introduced EVDO service in San Diego and Washington, D.C. EVDO offers speeds of 300 to 500Kbps, which is faster than existing GPRS (10 to 50Kbps) and 1xRTT (60 to 70Kbps). Verizon Wireless will extend EVDO service to additional markets over the next several years. At CTIA, LG Electronics was showing off its LG8000 mobile phone, which includes a megapixel camera and EVDO support. Prices and carriers for these upcoming phones were not available at press time. ▲

Nokia's 7610 phone includes a unique design and a 1.2-megapixel digital camera built in.

A Mobile TLD?

Is Google's mobile address www.google.com/mobile or mobile.google.com? If you're not sure, you may be interested by the new TLD (top-level domain) a group of nine big-name wireless players are planning. In March, Nokia, Vodafone, Microsoft, HP, Hutchinson 3G, Samsung, Sun, Orange, and the GSM Association put together a bid to jointly manage a TLD specifically for mobile content. In addition to making it easier to find mobile Web sites, a mobile TLD could provide a more consistent look and feel for mobile users. The group submitted its proposal to ICANN, which now must decide whether to allow the new TLD. Provided it approves it, ICANN would also bestow the TLD with its official name. ▲

(By the way, it's mobile.google.com.)

Wireless At The Ballpark

SBC Communications is doing more than just giving its name to the ballpark the San Francisco Giants call home. SBC has installed 121 802.11b wireless access points throughout the stadium so fans can get online from just about anywhere.

The Giants are planning a Digital Dugout full of extras for Wi-Fi users, including access to real-time stats, scores from other games, and, of course, the ability to order merchandise online. In the suites, fans will have access to an HP desktop computer with a Wi-Fi connection and an expanded Digital Dugout where they can watch archived video footage and even order food.

SBC plans to provide wireless access free throughout the 2004 season. In 2005, it may begin charging for Internet access. Fees for SBC's FreedomLink are $7.95 a day or $19.95 a month.

In addition to flooding SBC Park with Wi-Fi waves, SBC is also planning to install FreedomLink in UPS and Mail Boxes Etc. stores. SBC plans to install wireless access in 1,500 stores by the end of this year, and installation will continue into 2005. ▲

Leave your glove at home and take an 802.11b compatible device, such as this HP iPAQ 5500, to the San Francisco Giants' SBC Park, instead. SBC is providing free wireless Internet access to users throughout the 2004 season.

MVNOs Come To North America

The big wireless carriers in North America own and operate their own networks. That may change, however, as many companies are realizing they don't need the expense of a network to provide wireless service. Rather than spend millions on network infrastructure, MVNOs (mobile virtual network operators) are buying service from wireless carriers and reselling the service to consumers. Most MVNOs offer either a brand name or additional features to their wireless service to provide customers with more value. Virgin, one of the largest MVNOs in North America, uses its brand name to sell service to a fairly young demographic. An agreement with Sprint PCS lets its traffic flow across the Sprint PCS network.

Qwest recently announced it has reached an agreement with Sprint PCS and will offer added features to its wireline customers when they add Qwest wireless service. For example, customers will be able to share a single phone number between their wireless and wireline phones. Some customers may even receive free night and weekend long distance on their home phone.

There's also talk that AT&T Wireless may make a comeback to enter the fray. The money Cingular paid for AT&T Wireless didn't buy it the AT&T name, which reverts back to one-time parent company AT&T. There's speculation that AT&T could launch an MVNO service in the near future. ▲

More Tech Starting With "Wi" & "802"

In technology years, 802.11g is on the verge of a midlife crisis. That means we're almost due for a new 802.11-based wireless networking specification. Not to fear. The good engineers of the 802.11 working group are working on a new and improved wireless standard dubbed 802.11n. Work is still in the early stages, so it's unlikely that we'll see any 802.11n units before 2006. Currently, experts expect 802.11n to have speeds near 108Mbps, making it roughly twice as fast as 802.11a and 802.11g. It's not yet known if 802.11n will operate at 5GHz or 2.4GHz or if it will be backward compatible with existing Wi-Fi technology.

The problem with 802.11 is that it has a relatively small range. You need multiple access points to cover larger areas. The 802.16a, or WiMAX, specification defines wireless broadband delivery at distances up to 30 miles (point to point). 802.16a, however, defines fixed wireless access, making it less interesting to mobile enthusiasts. 802.16a revision e (commonly known as 802.16e), however, adds support for mobile networking. Some enthusiasts envision 802.16e phones, plus PDAs and notebooks, but as always we'll have to wait to see how things shake up.

Intel expects to have 802.16a chips on the market in the second half of 2004, but the 802.16a working group isn't expected to approve the final specification until next year. WiMAX-certified products aren't expected until mid-2005. ▲

At Your Leisure

Compiled by Samit G. Choudhuri & Chris Trumble

Plug In, Sit Back & Fire Away

The entertainment world, at least where it pertains to technology, morphs, twists, turns, and fires so fast it's hard to keep up. But that's exactly why we love it. For the lowdown on the latest in PC entertainment, DVDs, consoles, and just stuff we love, read on.

Unreal Tournament 2004

Third Time's A Charm

Unreal Tournament was released to much praise in winter 1999 and was followed by the much-anticipated Unreal Tournament 2003 in fall 2002. The latter rested on its laurels and just didn't make the same type of impact as the original. (Removing the Assault mode didn't help UT2003's cause.) The game just wasn't relevant when you could play Battlefield 1942 instead. Developer Epic Games picked up on this, updated the game, executed a top-notch plan, and came up with a superb multiplayer action game in the form of Unreal Tournament 2004.

Graphics are noticeably better than UT2003, even though UT2004 is based on the same engine. It may not be Far Cry (reviewed last month), but you won't be complaining about the graphics for some time. As long as you meet the minimum requirements, you will be able to enjoy the game because it scales down from high-end PCs very nicely. Music and audio effects are what you'd expect, and the game now also supports voice features for the extra personal touch.

There are 10 modes of play (look them up on the Web site), but we'd happily have paid $39 (there's a bundled $10 rebate for UT2003 owners, so upgraders can get the game for $29) for just the Onslaught and Assault (new and improved since the original UT) modes. Team Deathmatch, Deathmatch, Capture The Flag, and the others are just icing on the cake. Onslaught uses elements of Battlefield 1942 and Tribes. Two teams go against each other with the goal of destroying the enemy's power core. Thing is, you have to destroy and capture all the nodes in the field, which attach to the core before you can destroy the core. This adds a major strategic element to the action. Overall, the game design has been very well thought out, and the addition of land and air vehicles makes that mayhem even more enjoyable. Download the demo from the official Web site and try it out.

The game does feature good single-player mode, but it's really more of a training ground to improve your skills before you head online. The bots in the single-player mode continue to excel, and at times you feel like you're playing real players. This is not a game to buy for playing solo. If you're a multiplayer action fan, you owe it to yourself to buy UT2004. ▲

ESRB: (M)ature • Atari

$39.95 • PC

www.unrealtournament.com

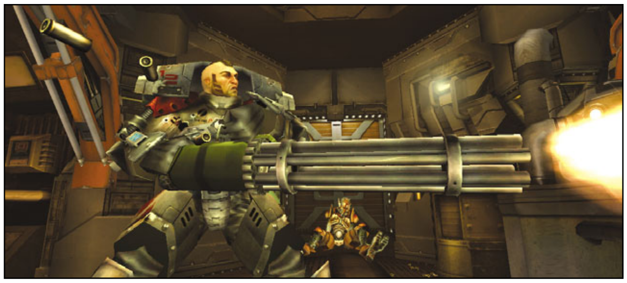

Is it really any surprise you don't get asked out much?

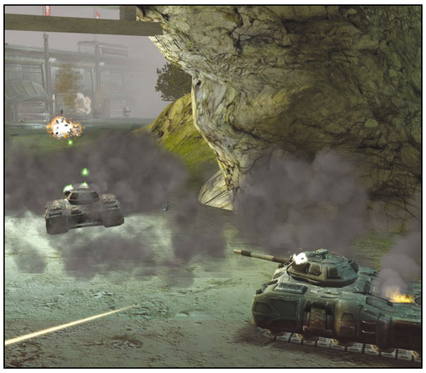

The addition of land and air vehicles in the Onslaught mode does not create imbalances during competition. It's all oh so good.

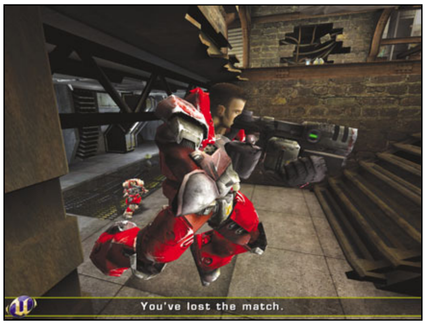

Win Team Deathmatch and you can look triumphant while the rest of us look on.

Metal Gear Solid: The Twin Snakes

Same Song, Second (Prettier) Verse

In case you haven't played the original PS version of this third-person action game, critics and many fans alike maintain to this day that MGS was the best game ever released for the PlayStation. That's high praise when you consider its competition, but the game was incredible. It was in many ways a pioneer of the current stealth action craze that has given rise to hits such as Splinter Cell and a host of others.

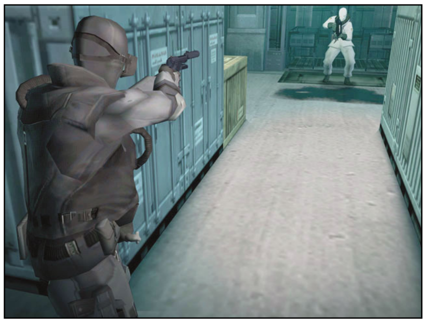
This scene look familiar? Eat your heart out, Sam Fisher; Solid Snake is the original stealth warrior.

The Twin Snakes doesn't offer much in the way of new content, but it does take a very good game to new heights. Improved in-game graphics, snippets of new footage in the game's already great cut scenes, and slightly reworked play mechanics that incorporate elements from its sequel (Metal Gear Solid 2: Sons of Liberty), make MGS:TTS a blast to play even if you played the 32-bit original to death.

And if you never experienced MGS the first time around, take our word on this one: Run, don't walk, to your favorite game retailer and pick this one up.

ESRB: (M)ature • Konami

$39.99 • NGC
www.konami.com/twinsnakes

Tom Clancy's Splinter Cell Pandora Tomorrow

Get Silent & Deadly In Multiplayer

The Xbox version of Splinter Cell debuted for the holiday 2002 shopping season and was promptly followed by the even more attractive PC version in early 2003. On the heels of those games comes Pandora Tomorrow for PC and Xbox. Little has changed: You once again play in the stealthy shoes of super-secret agent Sam Fisher as you try to thwart bioterrorism. It's basically a different campaign with improvements to the interface. Graphics are based on the previous engine but are crisper and show more detail. The in-game audio is highly effective; *listening* will often keep you alive. The PC version looks pretty amazing; Xbox gamers will not be disappointed.

The biggest improvement to Pandora Tomorrow is the multiplayer mode, without which the game would be considered an overpriced expansion pack (which it isn't). Up to four players can play at any time in a variety of combinations—all focused on stealth between spies (like Sam Fisher in the third-person) and spy-hunting mercenaries (in first-person). Teamwork comes in handy, and stealth kills are a thrill. The multiplayer mode will keep you coming back after you finish the single-player mode.

Sam Fisher eats a solid snake for breakfast every day. Yummy. It does the body good.

ESRB: (T)een • Ubisoft Entertainment

$39.99 Xbox • 49.99 PC
www.splintercell.com

The Suffering

Doin' Hard Time

Survival horror games are pretty common these days, but few are especially scary, relying instead on extreme gore and the occasional startle to keep players interested. The Suffering offers B-movie splatter (that and the language more than explain the game's M rating) and plenty of jump-in-your-seat startles, as well, but you also get the dark, atmospheric creepiness that you'll find in games such as Konami's Silent Hill series. Variety is a bit lacking, as the outlandish monsters start to seem very familiar after a while, but by and large this game is a pretty good time for survival horror fans.

When guards fight side by side with armed inmates, you know things have gotten pretty bad.

For one thing, there are few places more unsettling than a maximum-security prison. The Suffering takes this concept and runs with it by placing your character, Torque, in a horrific penal institution on a secluded island that is the epicenter of a whole lotta bad. Shortly after Torque's arrival on death row (it's a long story), a mysterious, unnatural event occurs that leaves Torque fighting for his life against those who have fallen victim to the island's cruelty in the past. ▲

ESRB: (M)ature • Midway

$49.99 • PS2, Xbox
www.midway.com

PS2 HDD/Final Fantasy XI
Sony Ups The Ante To 40GB

Because the PS2 is only the second console to get a hard drive, the obvious standard by comparison is Microsoft's Xbox. At the end of the day, Sony's offering is better in some ways and not quite as good in others.

On the plus side, the Sony HDD is very easy to install and provides 40GB of drive space, which is a vast improvement over the Xbox's 8GB drive and is more than enough room for downloaded content and game save files. Although you can't save just any game to the HDD from an in-game menu (at present, only Final Fantasy XI lets you save directly to the HDD), you can easily transfer saves from your PS2 memory cards to the drive from within the PS2's browser interface. Creating and naming folders on the HDD is easy, too, so organizing your game data is a snap.

In addition, developers are already beginning to take advantage of the HDD to optimize system performance. Capcom's Resident Evil Outbreak, for example, lets you save a big chunk of files to the HDD to cut down on in-game load times, and in our tests, the time savings were pretty substantial. Because load times are one area where the Xbox has traditionally trumped the PS2, smart PS2 developers will capitalize on this feature moving forward.

The obvious downside is the HDD's price. The drive costs $99, and although it comes with a game that would sell for close to $50 on its own, it still pushes the total cost of the PS2 to roughly $319, including the Network Adaptor ($40; required to use the HDD) and assuming you bought your PS2 after the most recent price drop to $179.99. Still, the drive is a very cool gadget and, if Sony plays its cards right, could do great things for the PS2's online capabilities, which brings us to Final Fantasy XI.

We reviewed the PC version of FFXI in the March *CPU*, and we like the PS2 version every bit as much. Naturally, the game doesn't look quite as nice as the PC version on a rig with a high-end video card; we noticed a considerable amount of draw-in and occasional frame rate hitches as our PS2 worked to render the picturesque land of Vana'diel. From the perspective of a console gamer, though, the PS2 version's control scheme is much cleaner and more intuitive than using a keyboard and mouse—right up until you try to chat with fellow players, that is. If you've ever tried chatting with a software keyboard, you know it's an exercise in frustration, but adding a USB keyboard costs just a few bucks more and solves this lone control issue handily. And like the PC version of FFXI, the PS2 version installs fully to the HDD and lets you dispense with loading game discs to play.

FFXI for the PS2 is a great game, and the ability to finally play online with PC gamers makes it very interesting from a cultural and technological standpoint. As such, we think forward-thinking gamers (especially RPG fans) will find the HDD/Final Fantasy XI bundle a worthwhile investment. If other third-party developers get onboard and support the HDD, it may well be the best $100 most PS2 gamers have spent since purchasing their systems. ▲

Sony's 40GB HDD sells for $99 and includes the Square Enix MMORPG Final Fantasy XI.

ESRB: (T)een (Final Fantasy XI) • SCEA

$99 • PS2

us.playstation.com

DVD Byte by Todd Doogan

THE DIGITAL BITS
we know dvd
www.thedigitalbits.com

I probably shouldn't waste my space with pleasantries because May is a huge month for DVDs. Let's take a look: On May 4 we get a whole lot of TV with **The Bernie Mac Show: The Complete First Season**, **Gilmore Girls: The Complete First Season**, **Law & Order: The Second Year**, **The Osbournes: 2 1/2**, **Party Of Five: The Complete First Season**, **Sonny & Cher: The Ultimate Collection**, and **Three's Company: Season Two**. It's also a big samurai day: Look for Cartoon Network's **Samurai Jack: Season 1**, **Criterion's Samurai Trilogy** (a brilliant series of films), and Tom Cruise in **The Last Samurai**. Horror fans can rejoice on the 11th, with re-releases of Universal's classic horror films in Legacy Collection form (complete with minibusts of your favorite characters). Buy up **Dracula**, **Frankenstein**, and **The Wolfman**. Be on the hunt for stellar documentaries, as well, such as Oscar winner **The Fog Of War** and **Easy Riders, Raging Bulls**. Rounding out an already stellar DVD day, pick up animation hits **Invader Zim Volume 1: Doom, Doom, Doom**, **The Jetsons: The Complete First Season**, and **Jonny Quest: The Complete First Season**. The 18th is Special Edition day with **Enter The Dragon**; **The Good, The Bad And The Ugly**; **Paycheck**; **Wyatt Earp**; and **You Got Served**. Fans of Disney can fill their library with the new Treasures releases: **Mickey Mouse . . . In Living Color**, **On The Front Lines**, **The Chronological Donald Volume 1**, and **Tomorrowland**. Finally, the 25th brings us Bruce Campbell as Elvis in **Bubba Ho-Tep**, **Takashi Miike's Full Metal Yakuza**, **Akira Kurosawa's Stray Dog**, the unrated extended cut of **Underworld**, and a little independent film from New Zealand entitled **The Lord Of The Rings: The Return Of The King**. Like I said, save your money. May is a big month. ▲

Hot Shots: The Beauty Of The Game

Yeah, we know it's all about the gameplay. Sure, there are those who would have you believe graphics are relatively unimportant in the greater scope of things, but if you read *CPU* mag, you probably already know those folks are off their collective rockers. We want great gameplay combined with stunning graphics, and here are two upcoming games that show promise. Watch for them.

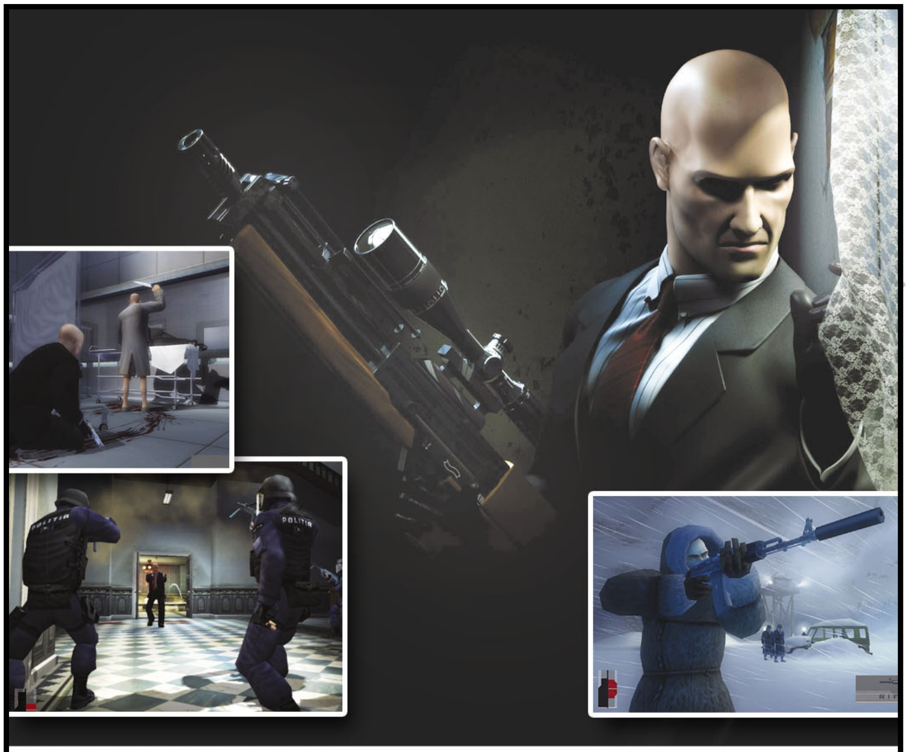

Hitman: Contracts (PC, Xbox, PS2). It's tough being a paid assasin, but you do what you're best at. Unlike Splinter Cell, you get to mix it up with stealth and action. Graphics have improved, and the developers have attempted to make the game more newbie-friendly (thank you!) by adding a learning curve. Learn more about it at www.hitmancontracts.com.

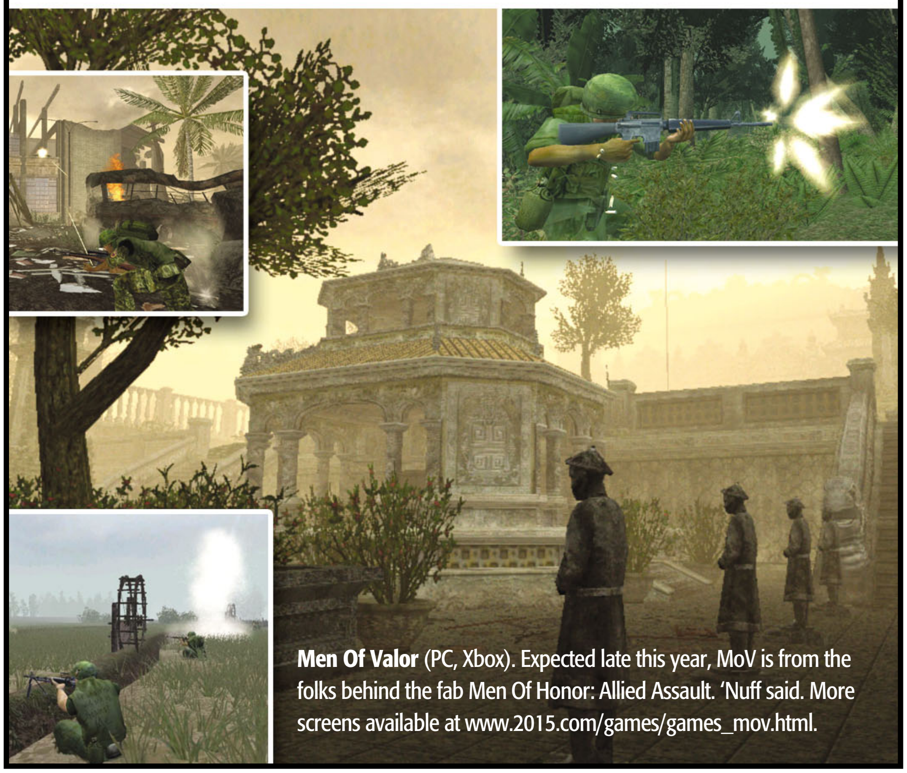

Men Of Valor (PC, Xbox). Expected late this year, MoV is from the folks behind the fab Men Of Honor: Allied Assault. 'Nuff said. More screens available at www.2015.com/games/games_mov.html.

Infinite Loop

Plant Stoplights

Totaling 100 million around the world with 40,000 new ones deployed every week, land mines cause 25,000 injuries and deaths each year. Aresa Biodetection, a Danish biotech firm, has genetically engineered a plant to detect land mines and react by changing color. The thale cress plant, a member of the mustard family, is usually a vivid green, but it will turn bright red when its roots detect traces of explosives in the soil. The scientists have altered the plant's usual reaction to stresses such as drought or cold temperatures so it only changes color after exposure to nitrogen dioxide, signaling the presence of land mines.

The question is, who sows the seeds in the existing minefields? The plants would be sown by clearing strips using conventional methods first and then planting them or by using crop dusting planes. No gardening tools required. ▲

Source: Wired.com

© 2003 PriceGrabber.com. Products/trademarks shown are not in any way intended to indicate an affiliation or endorsement of PriceGrabber.com by their respective holders.

With hundreds of digital cameras to compare, you'll definitely find one that clicks.

If you're into digital photos or video, finding the right camera is no longer a shot in the dark. That's because PriceGrabber.com provides the fastest, most convenient way to locate and research the best products and prices online. Simply choose your criteria, click, and within moments you're comparing cameras by brand, features, popularity and price. While you're at it, get the lowdown on products and merchants from buyers just like you. Bottom line? Whether you're shopping for cameras, computers, TVs, or thousands of other popular items, PriceGrabber provides all the information you need to make the best buying decisions.

• Computers • Software • **Electronics** • Video Games • Movies • Music • Books • Toys • Office Products

PriceGrabber.com
Comparison Shopping Beyond Compare.™

SOFTWARE TIPS & PROJECTS

Slipstream Your Next Windows XP Installation

IF YOU HAVE BEEN LIVING WITH WINDOWS XP SINCE THE OS LAUNCHED A COUPLE OF YEARS AGO, THEN YOU'VE ALREADY FELT THE PAIN OF A REINSTALL. SINCE THE ARRIVAL OF the XP Service Pack (SP1) last year, reinstalling the OS has become a two-part process of installing the original XP disc and then downloading and installing the SP1. There is an analgesic for your reinstallation pain, however; it's called slipstreaming. Developed as a way to streamline network installations of WinXP, slipstreaming lets you apply the SP1 (now SP1a) update to the original XP OS files before you do a reinstallation, saving you from a long and tortuous extra step.

Slipstreaming requires that you make a bootable CD with the updated XP system files, so you will need a CD-R\RW drive and good burning software. Note that by the time you read this, Microsoft may have released Service Pack 2 for WinXP. Users of the beta report that you can use the slipstreaming method outlined below with SP2, as well. Also, this technique works on both WinXP Home and Professional. The same SP1 update can be applied to both installations.

Preliminary Preparations

You'll need to create three new folders: one to hold the compressed SP1 files (such as C:\SP1down), one to hold the extracted SP1 files (C:\SP1), and one to hold the XP system files (C:\XPdisc). Because of the extracted and temporary files required for CD burning operations in this project, count on needing several gigabytes of free hard drive space.

The slipstream function is only available in the network version of the XP Service Pack, so you need to go to www.microsoft.com/windowsxp/pro/downloads/servicepacks/sp1/network.asp for the link to the English language, network version of SP1a, which you should click and direct to download into the new SP1down folder you made. It takes up about 125MB.

Now we need to copy the full contents of your original WinXP disc to its temporary folder. Insert the WinXP disc and exit from any Welcome screen that runs automatically, then right-click the CD drive in My Computer and click Explore to open its contents. Simply highlight and click and drag the entire contents of the root directory of the CD to the XPdisc folder on your hard drive.

The next step is to extract the SP1a files. Click Start, Run and type **command** in the Open text box. Navigate to the subdirectory holding the compressed SP1a pack by typing the old change directory command: **cd\c:\sp1down**. Type **dir** at the command prompt so you can see the actual name of the compressed file. In our case it was Xpsp1a_en_x86.exe. At the command prompt, type ***file name* –x**, which will prompt you for a target folder. Either type in your target (**c:\sp1**) or use the Browse button to navigate there and then click OK to extract the files.

Still using the command line interface, change to the /update subfolder in the target SP1 folder by typing **cd\sp1\update**. In order to slipstream the service pack into the WinXP files,

WinXP Tip Of The Month

Batch Renaming

If you use a digital camera, you probably have piles of images with nondescriptive file names such as IMG_1.jpg. Rather than rename each of these files individually, let WinXP help you quickly and easily rename the whole batch at once. In the folder that contains the images you want to rename, use the View menu to switch to Thumbnail or List mode. Use the SHIFT key to help you select the group of files you want to rename.

Now right-click the last file in the group and choose Rename. Name the file something you want to use for the entire sequence, such as Summertrip.jpg. Now press ENTER, and the entire run of files will be renamed in numbered sequence: Summertrip (1).jpg, Summertrip (2).jpg, etc. The only glitch in this trick is that the last file in the group will not have a parenthetical number, but you can right-click that file easily enough to add the final number in the sequence. ▲

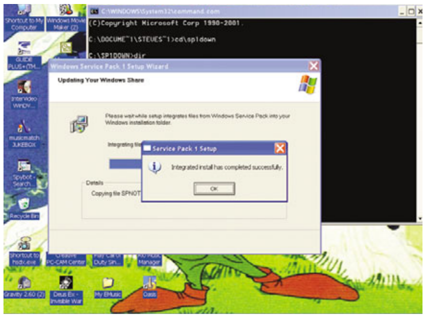
C:\WINDOWS\System32\command.com
Windows Service Pack 1 Setup Wizard
Updating Your Windows Share
Service Pack 1 Setup
OK
Details
Shortcut to My Computer
Windows Movie Maker (2)
Recycle Bin

```
(C)Copyright Microsoft Corp 1990-2001.

C:\DOCUME~1\STEVES~1>cd\sp1down

C:\SP1DOWN>dir
```

Please wait while setup integrates files from Windows Service Pack into your Windows installation folder.

Integrated install has completed successfully.

Copying file SPNOT

The slipstream function in SP1 lets you integrate the SP1 revisions into Windows XP's original installation files.

type **update /s:***path to XP disc folder* (where the italicized text is the directory containing the files we copied from the original XP installation disc, such as **update /s:c:\xpdisc**). You should now see a new Service Pack 1 Setup Wizard window that will show you the progress of the SP1 files integrating with the OS. When this is complete, you can close the command prompt window.

Boot It Up

At this point, you could simply insert a blank CD-R and drag the files from your XPdisc folder to the drive to make a nonbootable Windows installation. You would need to boot with a floppy (with CD-ROM drivers enabled) and type **setup** at the command line of the CD drive to start the installation.

But if you want an even more streamlined process, make the slipstreamed CD bootable. To do this, get the handy freeware program (donations accepted) IsoBuster at www.smart-projects.net/isobuster. This program lets you copy the boot image from a disc. Install the program and then reinsert your original Windows XP installation disc. Start IsoBuster and direct the program to view the CD drive with your original WinXP disc. In the left pane of IsoBuster, highlight the Bootable CD label in the directory tree and then right-click Microsoft Corporation.img in the right pane. This is your boot image. Use the Extract command and direct the dialog box to save the boot image to your XPdisc folder.

Now to burn it all. We used Ahead Nero, but make sure you have a recent version. We used 5.5.10. Press CTRL-N to start a new compilation and choose CD ROM [Boot] from the left menu. In the Boot tab, click the Image File radio button and use the Browse button to find the Microsoft Corporation.img file in your XPdisc folder. In the file browser, we had to choose All File from the bottom drop-down menu to ensure that the IMG file was visible. Click Open.

Precision Burning

Your settings must be correct in Nero for the disc to boot properly. In the same Boot tab, make sure the Enable Expert Settings checkbox is selected and that No Emulation is chosen in the Kind Of Emulation drop-down menu. By default, the Load Segment Of Sectors box should read 07C0. In the final box, Number Of Loaded Sectors, change the default 1 to 4. Now go to the ISO tab and make sure that the following options are set using the checkboxes and radio buttons.

File/Directory name length: ISO Level 2

Format: Mode 1

Character Set: ISO 9660

Joliet: selected

Relax ISO Restrictions: all three boxes selected

You can use the Label tab to give your disc a unique volume label. Many people use the original labels used for the various types of WinXP installation discs. For example, WXHCCP_EN is the label for non-OEM versions of WinXP Home. We've had no trouble making bootable slipstream discs with custom labels. Click New.

Now you can highlight and drag all of the files in your XPdisc folder into Nero's main file window. Choose Write CD from the File menu to open the Write dialog box. Click Write to make your copy. You should now have a bootable CD that will load Windows XP with SP1 already integrated.

For this to work, you need to have a PC whose BIOS lets you boot from a CD. Also, because the bootable slipstream disc probably uses between 500MB and 550MB of your disc, there is room to add a subdirectory of drivers and utility programs you will want to reload onto your system when you reinstall the OS. Finally, remember that making a bootable, slipstreamed WinXP installation disc doesn't skirt the copy protection that Microsoft has put in place. You will need to keep handy the original disc and its product key number to enter at the beginning of installation, and the activation process still does its check to see that the OS is only being reinstalled to the same machine.

Thanks to both TheElderGeek.com and Windows-Help.net, whose respective walkthroughs of this procedure we used as reference throughout.

by Steve Smith

Registry Tip

The New Window Trick

To avoid screen clutter, most of us set the Explorer Folder Options so that double-clicking a folder opens its contents in the same window. In some cases, however, opening a folder in a separate window is preferable; this Registry hack lets you add the option to your context menu. In the Registry Editor, go to HKEY_LOCAL_MACHINE\SOFTWARE\CLASSES\DIRECTORY\SHELL and create a new key named Opennew. In the right pane, double-click Default and type **Open New &Window** in the Value Data box. Back in the left pane, create a new key under Opennew and name it Command. Set the new key's default value to explorer %1, click OK, and close Regedit. This change requires a system reboot to take effect; it adds an Open New Window command when you right-click a folder. ▲

Infinite Loop

What's up, Doc?

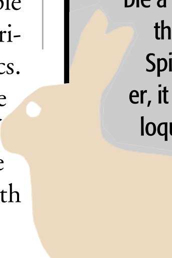

According to NASA.gov, the recently discovered evidence of water in Mars' Meridiani Planum region by Mars Exploration Rover Spirit lends strong evidence to the theory that conditions on the planet may once have been hospitable to life. This puts an interesting twist on earlier images captured by Spirit's twin, Opportunity. NASA scientists speculate that Opportunity's images, which resemble a rabbit, are actually parts from the rover or its lander. In light of Spirit's recent discoveries, however, it seems possible that images of a loquacious black duck or a short humanoid resembling a Roman centurion are forthcoming.

WARM UP TO LINUX

Sorting Out Regular Expressions

REGULAR EXPRESSIONS, OFTEN SHORTENED TO JUST REGEXP, ARE SYMBOLS THAT ARE USED TO MATCH PATTERNS IN TEXT. A SIMPLE EXAMPLE OF A regular expression most people are familiar with in computing is the asterisk (*), which matches one or more of any grouping of characters. Thus, a*c could match abc, abbc, abac, abdefghic, and so on. However, matching needs usually aren't as simple as this. If you use or plan to use such tools as sed, awk, or Perl, you may have to write giant links of characters to tell the computer how to recognize ranges of text you want it to look for, such as this example: ^<(?:\d\d?\.){4,5}\d{14}\.[a-f0-9]{8}\@\S+>$

So, just what does this expression mean? To help us figure it out we can use the Regex Coach graphical tool. (It's important to note that there's more than one type of regular expression. For our purposes, we're focusing on Perl regular expressions, which many different languages and applications use.)

Get Regex Coach

To get Regex Coach, go to weitz.de/regex-coach. On the main page, click the Download And Installation link and click the Linux version (a Windows version is also available) to download the file to your download directory. After downloading it, uncompress and unpack it by typing **tar xzvf** ***file name***. (For example, if the file name were regex-coach.tgz, you'd type **tar xzvf regex-coach.tgz**.) Use the cd command to change to the new regex-coach directory, and run the program by typing **./regex-coach &** at a command prompt. Adding the ampersand lets you run the program, plus free up your command prompt for other tasks.

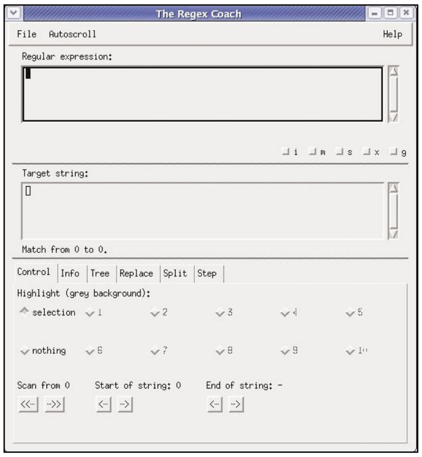

Match from 0 to 0.

Scan from 0

Start of string: 0

End of string: -

This image shows the Regex Coach tool running under Linux. A Windows version of Regex Coach is also available.

We will use the coach tool to walk through an expression. We will start by inputting the regular expression example we listed earlier into the coach's Regular Expression text area. Now we'll click the tool's Info tab. You'll notice that there's an exclamation point at the end of every line, but don't let this confuse you. The exclamation points aren't part of the expression.

As the Info tab explains, this expression begins with a start anchor, which looks like this: ^ (The start anchor specifies that what follows should be at the beginning of the line.) The Info tab then points out that after the anchor there must be a less than sign (<), so the line has to begin with a < character. After the less than sign, there must be (?:\d\d?\.){4,5}. These characters actually form a regular expression, but Regex Coach doesn't expand on what it signifies. Rather than try to decipher what this particular entry means, we'll open a second instance of Regex Coach from another command prompt and insert this smaller portion in the Regular Expression text area.

According to the Info Tab in the second instance of the application, (?:\d\d?\.){4,5} tells the program that is processing the expression to look for the expression \d\d?\. repeating four or five times in a row. We'll delete this expression in the second Regex Coach's text area and replace it with \d\d?\. This tells us that this entry expands to any numeric digit, followed by the expression \d? and a period. When we evaluate \d? we discover that this expression means no or one instance of any numeric digit. So, the text matching this expression so far might look something like this: <3.3.3.3.

Now we'll return to the first instance of Regex Coach. The next expression within our large jumble of characters is \d{14}. Again, we'll copy this into the second open coach program, which tells us that this piece is looking for a numeric digit listed 14 times in a row. So a potential match might be: <3.3.3.3.88888888888888

According to the Info tab, this part of the expression is looking for a period and the results of the expression [a-f0-9]{8}.

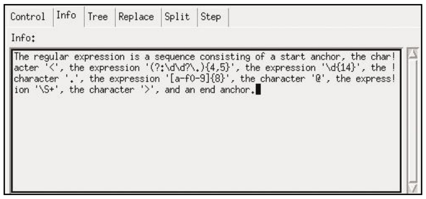

```
The regular expression is a sequence consisting of a start anchor, the char!
acter '<', the expression '(?:\d\d?\.){4,5}', the expression '\d{14}', the !
character '.', the expression '[a-f0-9]{8}', the character '@', the express!
ion '\S+', the character '>', and an end anchor.
```

This example shows Regex Coach's Info tab with our example expression explained (in mostly plain English).

Not Happy With Your Linux Distro?

For the übergeek, there's Linux From Scratch (www.linuxfrom-scratch.org). Go to the site, select the closest mirror, and you'll be sent to a site dedicated to helping people build their Linux installation from the ground up. Whether you want to design your own specialized distribution, get yourself some geek cred, or just learn what's under the hood on a really intimate level, LFS might be for you. Your best bet for getting started is to read the FAQ. ▲

This shorter expression within the larger one is looking for any one of the letters a through f and the digits 0 through 9 to be repeated eight times. So a match (including the period) might be: <3.3.3.3.88888888888888.dddddddd

We're still not done, however. Next, the expression is looking for an at (@) symbol and the results of \S+, which refers to any character that isn't a space appearing one or more times. Finally, there's a greater than symbol (>) at the end. This symbol must be the last thing on the line. So, the match might look like this: <3.3.3.3.88888888888888.dddddddd@***>

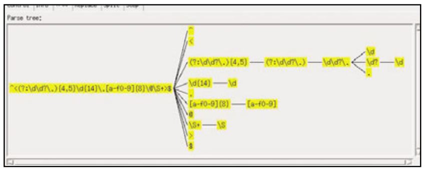

The Tree tab in Regex Coach shows our example expression displayed in graphical form.

More Fun With Expressions

Clicking the Tree tab in Regex Coach gives us a more visual look at the expression. Here, each piece of the expression is shown in a graphical format, with each smaller expression given its own branch. In some cases, subportions of each expression are broken off into their own branches. Think of the tree as the structure for how this entire regular expression is evaluated, not in one fell swoop but by determining the nesting and then working out each individual piece.

We'll click the Step tab to try another cool feature. In the upper instance of the box labeled Target String: type something that you think will match this regular expression. Typically, you would type an example of the text you are trying to catch with this regexp. We will use the example match from above: <3.3.3.3.88888888888888.dddddddd@***>

The same phrase appears in the program's Target String: area. Now, click the Start button at the bottom of the program. The opening letter will be highlighted, and the Regex Coach tells us that there's a submatch, meaning that so far the phrase matches our regular expression. Click Next Step and you'll see that there's another submatch, as more of the phrase is considered. As you continue clicking Next Step, you will eventually reach a point where the phrase no longer matches or the entire phrase matches. If your regular expression isn't a full match, the Next Step button goes away and you're told "No Match." In this case, you'll see something like "Match from 0 to 38."

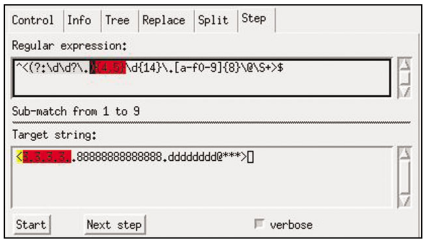

Regular expression:

```
^<(?:\d\d?\.){4,5}\d{14}\.[a-f0-9]{8}\@\S+>$
```

Sub-match from 1 to 9

Target string:

```
<3.3.3.3.88888888888888.dddddddd@***>
```

☐ verbose

The Step tab in Regex Coach shows our example expression and test match phrase in place, with the opening marked as matching.

Even better, you'll know immediately if it's a match because when you enter your value in the Target String text box in an attempt to match your regular expression, a match turns yellow.

Give Regex A Spin

The program includes other features, which you can play with to learn more about. The regular expression we've been working with has something to do with message identification strings in email headers; it's used in an antispam program to try to identify mail you don't want to receive.

To learn more about Perl regular expressions, check out www.perl.org, and read the appropriate documentation. You can also type **man perlre** at a command prompt if you have Perl installed. As of this writing, you would want to go to www.perldoc.com/perl5.8.0/pod/perlre.html for the latest version of the documentation.

There are other types of regular expressions. For example, if a tool tells you that it uses POSIX regexps, you can type **man 7 regex** at a command prompt to learn the rules for these. Other types include basic and extended regular expressions, both of which you can find more about by typing **man grep** at a command prompt and scrolling down to the section beginning with REGULAR EXPRESSIONS. CPU

by Dee-Ann LeBlanc

Infinite Loop

Hand Over The Cash

Not only do left-handed people have a higher tendency of being accident-prone and/or insane, they are also statistically more likely to be geniuses. Case in point: Four of the five original designers of the Macintosh were left-handed, as is Bill Gates. Being a southpaw has paid off for Gates: If he put all his money under his mattress and fell out of bed, he would have 18 minutes to prepare for the landing. If he added all the interest on his money, he might never land.

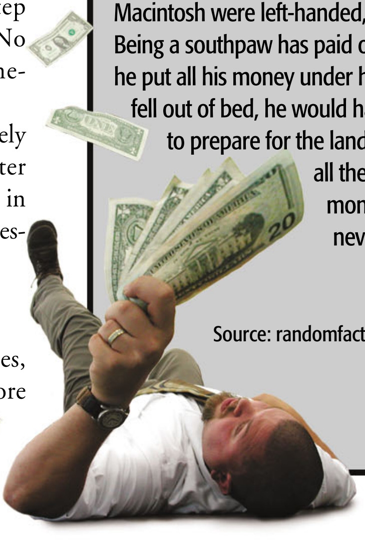

Source: randomfacts.diaryland.com

Hearing is Believing

PC WORLD
WORLD CLASS
2003

revolution 7.1

High-Definition Surround Sound for Your PC

Revolution 7.1
- PCI solution for high-definition surround up to 7.1
- 24-bit/192kHz output up to 8 channels
- Professional 24-bit/96kHz recording
- Real-time stereo-to-surround conversion
- Incredible 107dB SNR provides pristine audio

Sonica Theater
- USB solution for surround up to 7.1
- Up to 24-bit/96kHz audio output
- Professional 24-bit/96kHz recording
- Real-time stereo-to-surround conversion

TRUE 24 bit 192 kHz High Definition

M-Audio's audio cards deliver immersive surround sound on your computer, rivaling that of the best movie theaters. Games, movies, and music—even your stereo CDs and MP3s—all sound more alive and intense with M-Audio surround solutions. Now you can enjoy amazing new cutting-edge audio standards like 7.1 surround sound and high-definition 24-bit/192kHz fidelity. Support for popular surround formats including Dolby Digital and DTS, as well as gaming standards like Sensaura, EAX, DirectSound and A3D, insures that you'll hear everything exactly the way it was intended.

Only one name is found in more studios and is used by more music professionals than any other: M-Audio. Now we've distilled all that experience into the most powerful, high-fidelity PC surround sound solutions available. In fact, PC World recently gave the Revolution 7.1 its highest accolade–the coveted World Class Award 2003.

Get M-Audio. Get M-Powered.

available at
COMPUSA®
WHERE AMERICA BUYS TECHNOLOGY.℠

M-AUDIO
www.m-audio.com

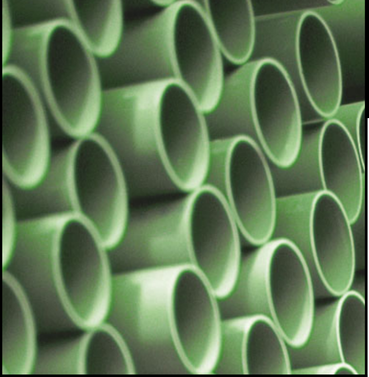

Shavings From The Rumour Mill

by Mike Magee

Intel's Numbers Games

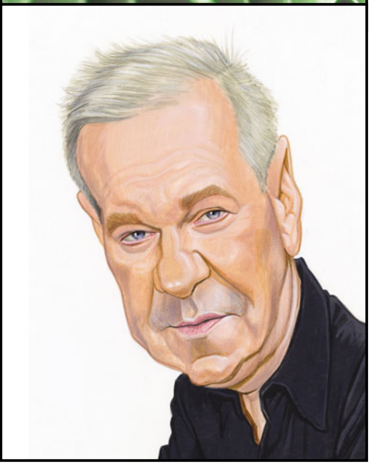

Mike Magee is an industry veteran. He cut his teeth on ancient products like the Dragon and the Japanese PC platforms long before the IBM-PC won. He worked for a corporate reseller in the mid-'80s and saw the Compaq 386 sandwich box and every GUI known to humankind. Mike decided that the way to go was the Interweb around 1994 after editing PC mags in the late '80s and '90s. A co-founder of The Register, Mike started the chip-driven INQUIRER (www.theinquirer.net) in 2001. He has contacts from top to bottom in the business, spanning the entire chain, who help him root out interesting rumours and speculation.

A lot's happened since I wrote my last column for *CPU*. Some of it is frankly unbelievable, but these rumours are about CPU facts, not science fiction. I first heard the rumour that Intel was going to use model numbers for its CPUs about four weeks before it happened. There's not much we need to know about the scheme. The 580 is a Pentium 4 that runs at 4GHz and that will be available during the fourth quarter. There's still some leeway on the Pentium 4 architecture, so the latest road maps we've seen are showing a 580+ for next year. Plus what?

Celerons, being a cut-down Pentium 4, use 3xx numbers and Intel tells us, now that the facts are out, that the higher the number, the more "featured" a chip is. This is quite hard to understand unless you think a microprocessor is like a car, rather than a carburetor. We all know that graphics cards, memory (not just on graphics cards but on a system), hard drive speeds and operating systems make a difference, too.

Some wags suggested last April Fool's Day that AMD had followed Intel's lead, but so far this has proved to be a jape. It can't be that long though until AMD does re-engineer its "PR rating" plans. Intel's scheme starts, oh, about now. It's already started trying to get us all to think 580 rather than 4GHz.

> **We know that Intel will have a dual core Pentium M next year, but that's only because an engineer we talked to at a recent conference got tipsy on a Bud Lite and spilled the beans.**

We believe the Netburst branding that was the main plank of Intel's Pentium 4 family will disappear sometime next year. Future Intel desktop processors and mobile processors will hardly be any different: Both will ultimately be based on the Pentium M which, it's worth noting, Intel never sold on the premise of its megahertz but on its performance and other features. Like the desktop Pentium 4s and Celerons, future 90-nanometre Pentium Ms will have numbers too. By the time 2004 is over, there will hardly be a chip that doesn't have a number rather than a gigahertz tag.

This, I believe, is part of a tenderizing process Intel is engaged in, which in a couple of years' time will lead to you buying a PC with a dual or quad microprocessor, a high-speed bus, 64-32 extensions, an internal memory controller, and HyperThreading. What will its number be? We don't know yet. But moles whispered to us earlier this year it's already got a code name and that's Conroe.

Intel has a fierce competitor in AMD. While reviews and much of the news this year has been all about the Athlon 64, the Athlon FX, and the Opteron server chip, we mustn't forget that AMD has roadmaps and has already been working on its K9 CPU for at least the last 18 months.

While Intel's Andy Grove is famous for uttering the expression "only the paranoid survive," in practice AMD is far better at keeping things secret. So far all we're sure about is that during late 2005 AMD is likely to introduce K9 microprocessors and will concentrate on low power consumption while still trying to up the ante on clock speeds. That's unless AMD decides to change its mind and use a similar numbering system as Intel to make life easier for us all.

We know a little bit about these chips. Like Intel, AMD is working on multicore versions of future technology. The marketing people at AMD have code names to leak, sorry, enthrall us with, too. In the first half of 2005 we're being promised Oakville, a 90nm notebook chip. Later that year, AMD goes code name crazy with three server chips called Egypt, Italy, and Denmark; two desktop chips called Toledo and Palermo; and a further notebook chip called Trinidad. It's almost tempting me to go on my holidays.

All these chips will also use 90-nanometre silicon on insulation technology but we suspect AMD has some dual-core CPU surprises as well. We know that Intel will have a dual-core Pentium M next year, but that's only because an engineer we talked to at a recent conference got tipsy on a Bud Lite and spilled the beans. ■

Send rumours to "Mad Mike" Magee at Mike@cpumag.com.

Technically Speaking

An Interview With John Stroozas, Director Of Engineering For Crucial

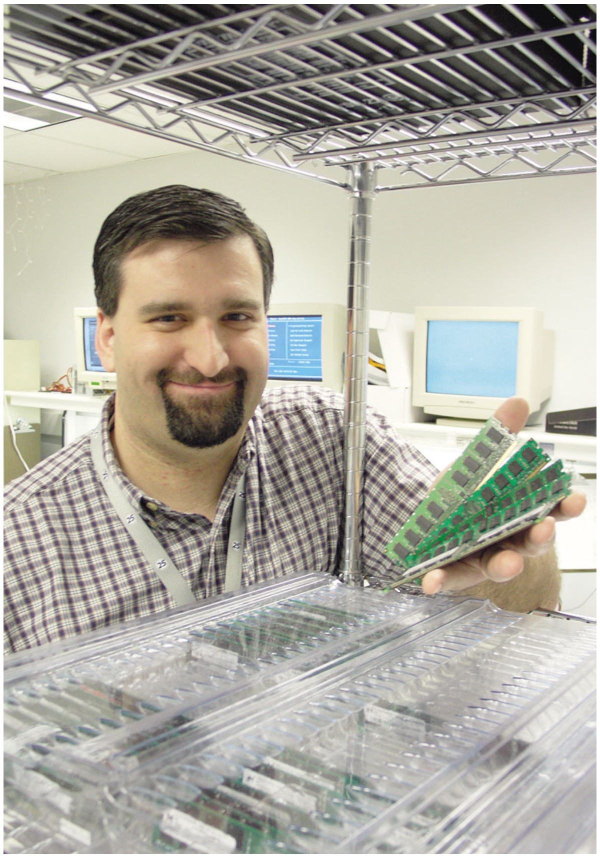

Everybody knows that DDR2 is big news for PCs—or at least it will be once speeds scale up a bit. What everybody doesn't get to know is the inside view of a technology change, such as DDR2, from the eyes of one of the lead engineers making it happen. John Stroozas is one of those geeks who got lucky. When he graduated college in 1997 at age 26, his first "real" job was working for Micron. Today, he is the director of engineering at Micron-owned Crucial Technology. He is Crucial's jack-of-all-trades on DDR2, doing everything from helping direct the company's designs to fielding tech support requests from end users. We figured that if anyone would be in a position to talk about DDR2 and the memory landscape at large, Stroozas would be the man. ▲

by William Van Winkle

CPU: DDR2 is the news item du jour, but most people aren't expecting major performance gains for a while. Should they care?

Stroozas: Basically they'll have to care when OEMs use it in their systems. It's really another one of those system memory changes that's going to be driven by the OEM market and chipset vendors as the new chipsets are released in the coming months that support DDR2. People will be required to use DDR2 with those motherboards and chipsets . . . As far as the enthusiasts go, they'll pick up on it because they're typically the ones that want to have the latest and greatest hardware out there.

CPU: And unlike USB 2.0 and 1.1, DDR2 slots are not compatible with today's DDR modules.

Stroozas: That is correct. It will not be backward-compatible. It will be a new socket that will specifically be used for DDR2.

CPU: Despite that the slots look very similar to one another?

Stroozas: Physical length is the same as a DDR1 module, but DDR2 has 240 pins. I'm talking about the regular or unbuffered or registered DIMMs having 240 pins on the module rather than DDR1's 184. You'll also notice that the center notch is placed a little differently. It's along the lines of the transition we had from EDO to SDRAM.

CPU: Starting out, what will be your DDR2 CAS latency?

Stroozas: CAS latency out of the chute for PC2-3200 will be 3. And then for PC2-4300, the CAS latency will be 4.

CPU: So, can you confirm some of the early buzz saying that, as with the Intel Northwood to Prescott change, we won't see much speed benefit from the initial iterations of DDR2?

Stroozas: That's hard to say right now since the chipsets and boards aren't really out in volume yet, and we actually don't have enough data to really show—

CPU: Aw, come on. You've played with this stuff in the lab.

Stroozas: Essentially, DDR2 400 has the same top theoretical bandwidth as DDR1 400. But that's just because it's essentially the stepping-stone for moving from DDR1 to DDR2, and it'll carry far beyond that bandwidth. Actually, DDR2 400 is expected to only have a very limited life span in application. Most of the mainstream volume of systems is expected to start at DDR2 533. That's where you should see the majority of the systems coming out. I'd equate that to what we saw with regular DDR with PC1600 and PC2100. There were a few 1600 systems out there, but the majority went to 2100 very quickly. Let me explain one thing about latency, too. A lot of questions initially have come through about the latency of DDR2 and why it's higher than DDR1, especially from enthusiasts. What happens is as you go to higher-on frequency speed, your latency is determined by your access time on the chip. We're continually trying to reduce the amount of time that it takes for that access to happen, but as clock speed kicks up, you have a shorter and shorter time in nanoseconds for each clock period. So, even though the access time may be the same or slightly smaller, if your clock frequency goes up, your CAS latency is actually going to appear a little higher even though the real-time in nanoseconds would be the same or smaller.

CPU: How well does DDR2 lend itself to overclocking?

Stroozas: Without really having a lot of data from real applications at this point, we're not positive, but it certainly seems that it will be similar to DDR1.

CPU: Is there a Moore's Law for DRAM?

Stroozas: Memory historically hasn't moved forward as quickly as processor speed, which has always created some issues and complaints from people. Well, with the migration from SDRAM to DDR, and now again from DDR1 to DDR2, you're seeing memory speeds increase significantly relative to their historical rate. We aren't necessarily going to catch up with the bus speeds of the CPU, but we're at least making leaps and bounds in our ability to keep pace and get the overall bandwidth increases that we didn't see in the past.

CPU: Will you be able to maintain that accelerated pace?

Stroozas: We're hoping so . . . and expecting to.

CPU: What enabled that acceleration?

Stroozas: The architecture used in the move from the EDO fast-page asynchronous days to the synchronous days of SDRAM and DDR have allowed for increased speeds, as well as the aggressive geometry shrinks in the industry. The smaller the geometries of the actual DRAM processes, the faster the parts are able to go.

Plus, there are other changes, especially with DDR1 to DDR2, which are going to make noticeable differences to the consumer. The obvious one is faster data rates and allowing parts to have a higher overall bandwidth for the system. But you're also talking about small geometries, which lead to lower power levels. You've got lower operating voltage on a DDR2 part. At release, DDR2 will be running 1.8 volts compared to 2.5 for DDR and 3.3 for SDR. Along with that reduced operating voltage, you also have a lower power consumption, which for mobile devices—laptops, PDAs, phones, anything that may adopt DDR2 parts—is a definite benefit because it will draw less power and take less battery power, yielding longer life for the part. You also get the benefit of generating less heat. Then, of course, the ever-migrating densities of DRAM are also keys. You're seeing with DDR2, Micron already has the 1Gb part validated at Intel, along with the 256 and 512 densities. The DDR2 spec will allow for densities up to 4Gbs currently.

CPU: Gigabit or gigabyte?

Stroozas: Gigabit at the IC level. At the module level, using 1Gb chips, you can get a 4GB module. Micron also has that validated at Intel, as well.

CPU: With lower voltages, there will be no problems migrating DDR2 into notebooks?

Stroozas: There are no problems expected. Desktops and servers will be the first applications you're going to see out there for DDR2, and expect that laptops will be shortly behind.

CPU: Looking to the future, you surely have access to the blueprints for next-gen technology. We mean stuff well beyond DDR3. Technologies like holography, DNA, molecular memory, and such. What strikes you as the most likely candidates for mainstream development?

Stroozas: That's a difficult question to answer. I imagine it will be similar to what we saw with the move from asynchronous memory to synchronous memory in that there were quite a few different types of memory out there being touted as 'the next memory,' but the further along we get with the maturity of DDR2, the clearer that picture will get.

CPU: For a young exec, you're awfully good at being politic and not giving a personal opinion.

Stroozas: [laughs] This one I honestly don't know the answer to. I don't know which way it's going to go at this point, but I think you will see a tighter working relationship between the CPU and chipset guys and the memory manufacturers than in times past, just to make sure that they have a viable solution that'll be good for consumers.

CPU: But you read the journals. You have gut feelings. Forget that you're a Crucial man for a minute and just speculate.

Stroozas: There's a lot of press lately about magnetic RAM. Some of the other, more interesting ones include organic memory, where it's grown on a substrate. Magnet RAMs are really interesting technology. The ability to have the speed and the density of a DRAM with the nonvolatility of a flash product has a lot of appeal because then you're looking at much reduced need for power and an awesome ability for portability, as well. But the big question really is how can you make it manufacturable in a cost-effective manner? A lot of technologies out there that could be built right now can give you the density suitable for PC main memory, but you have to get it cost-effective in a commodity type product. That's really where the challenge is. CPU

Subscribers can see extra interview content at
www.cpumag.com/cpujun06/stroozas.

What's Everyone Talking About?

ComputerPowerUser.com

Take Advantage Of All The Extra Benefits Of Being A Paid Subscriber Today!

Here's what's available to everyone:

- The CPU Web Log. Each weekday, our editors scour the Web to bring you the best high-tech (and occasionally offbeat) information in an entertaining way.
- Forums to share (or show off) your thoughts and expertise. Great whether you know it all or want to know more.

Subscriber Tip Of The Month

We've now made it even easier to browse our articles! From the home page, click the Browse All Articles link within the Search All Articles area. You'll then see a grey search area that allows you to select a publication (or all) and then view either the articles from the latest issue or articles within a particular date range. All results will then appear in alphabetical order. If you know exactly what you're searching for, click the Go To Detailed Search link in the upper right to search for articles by category or word/phrase.

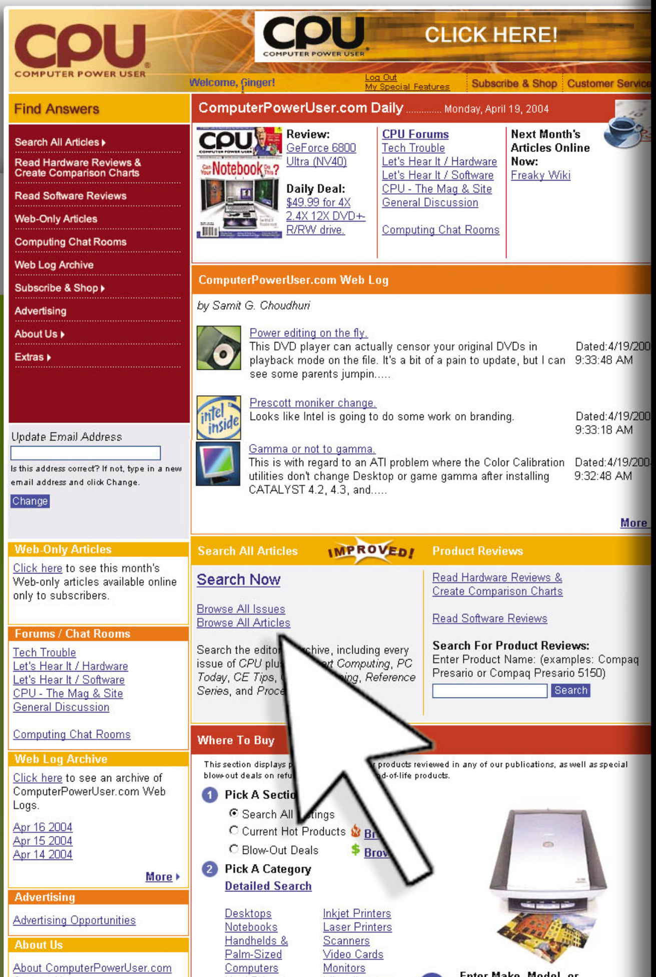

ComputerPowerUser.com Daily .... Monday, April 19, 2004

Review: GeForce 6800 Ultra (NV40)

Daily Deal: $49.99 for 4X 2.4X 12X DVD+-R/RW drive.

CPU Forums
Tech Trouble
Let's Hear It / Hardware
Let's Hear It / Software
CPU - The Mag & Site
General Discussion

Computing Chat Rooms

Next Month's Articles Online Now:
Freaky Wiki

ComputerPowerUser.com Web Log

by Samit G. Choudhuri

Power editing on the fly.
This DVD player can actually censor your original DVDs in playback mode on the file. It's a bit of a pain to update, but I can see some parents jumpin..... Dated:4/19/200 9:33:48 AM

Prescott moniker change.
Looks like Intel is going to do some work on branding. Dated:4/19/200 9:33:18 AM

Gamma or not to gamma.
This is with regard to an ATI problem where the Color Calibration utilities don't change Desktop or game gamma after installing CATALYST 4.2, 4.3, and..... Dated:4/19/200 9:32:48 AM

More

Update Email Address

Is this address correct? If not, type in a new email address and click Change.

Change

Web-Only Articles

Click here to see this month's Web-only articles available online only to subscribers.

Forums / Chat Rooms
Tech Trouble
Let's Hear It / Hardware
Let's Hear It / Software
CPU - The Mag & Site
General Discussion

Computing Chat Rooms

Web Log Archive

Click here to see an archive of ComputerPowerUser.com Web Logs.

Apr 16 2004
Apr 15 2004
Apr 14 2004

More ▸

Advertising
Advertising Opportunities

About Us
About ComputerPowerUser.com

Search All Articles IMPROVED!

Search Now

Browse All Issues
Browse All Articles

Search the edito chive, including every issue of *CPU* plu *rt Computing*, *PC Today*, *CE Tips*, *ing*, *Reference Series*, and *Proce*

Product Reviews

Read Hardware Reviews & Create Comparison Charts

Read Software Reviews

Search For Product Reviews:
Enter Product Name: (examples: Compaq Presario or Compaq Presario 5150)
Search

Where To Buy

This section displays p r products reviewed in any of our publications, as well as special blow-out deals on refu d-of-life products.

① Pick A Sectio
● Search All tings
○ Current Hot Products Br
○ Blow-Out Deals Brov

② Pick A Category
Detailed Search

Desktops
Notebooks
Handhelds & Palm-Sized Computers

Inkjet Printers
Laser Printers
Scanners
Video Cards
Monitors

Enter Make, Model, or

Only paid subscribers to *CPU* magazine get complete access to all these great features:

Latest Issues. Get a guided tour through the latest editions of our technology publications.

More hands-on hardware reviews. Are the latest products worthy of the marketing hype? We put products to the test, complete with comparison charts and links pointing to where you can buy the products online.

Even more great reading. Our Web site features every article from *CPU* and a searchable archive of more than 18,000 articles from our sister publications *Smart Computing*, *Learning Series*, *Guide Series*, and *Reference Series*. It's like getting five magazines for the price of one!

A showcase of articles from next month's issue. Impress your friends with all the insider information you get by reading *CPU*'s latest articles and product recommendations before they're available in print.

CPU
COMPUTER POWER USER

Under Development

A Peek At What's Brewing In The Laboratory

Doubly Dense

The highest capacity in the present generation of 3.5-inch hard drives is 80GB per platter. Many factors contribute to overall areal density, but one critical factor is magnetic isolation between tracks written onto the data storage layer in today's disks. Poor isolation means that tracks must be farther apart or that additional error checking is required, which impact capacity and drive speed.

Komag, the largest independent manufacturer of thin-film media for hard drives, just completed proof-of-concept testing of its DTR (discrete track recording) technology, which is supposed to yield roughly double the capacity of contemporary platters. DTR works just as well on today's longitudinal magnetic recording or next-generation perpendicular magnetic recording media. Komag expects DTR to facilitate platter capacities of 160GB to 240GB or greater—with 4TB per drive.

The key advance was being able to use NIL (nanoimprint lithography) to create a pattern of lands and grooves on the substrate prior to the deposition of the magnetic data layer. This allows for far greater precision in directing the head to the track's center and minimizing intertrack interference, which lets manufacturers loosen their recording head precision requirements. Additionally, because NIL technology can create surface features a mere 70nm in size, the nanoscale writing capabilities of NIL tools let disk manufacturers encode head precision information into the DTR disk track rather than relying on the costlier method of writing servo control instructions onto a continuous disk during drive assembly.

"The concept of patterning a disk's tracks has been recognized as potentially advantageous for more than 20 years," says Christopher Bajorek, Komag executive vice president, advanced technology. "But the problem has always been how one could achieve that in a cost-effective way. The advent of nanoimprint lithography, which was developed for general applications for future patterning of structures at a nanoscale level, when merged with a hard disk process, had the potential to accomplish this. We've successfully proved this principle, which has now moved into a prototyping phase."

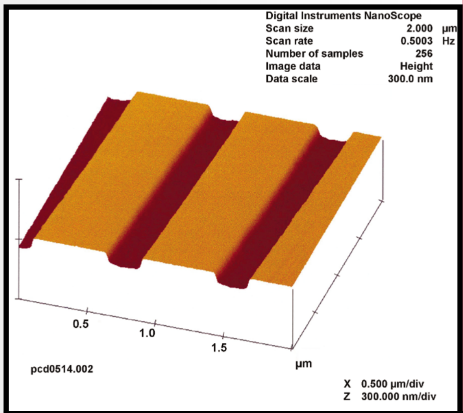

Komag is entering the prototyping phase for a nanoscale read/write technology that's called DTR that will double the data capacity of modern hard drives.

Part of Komag's achievement was developing NIL tools capable of creating DTR prototypes. Originally there were no NIL tools in the market. However, Komag is still cautious about delivering a timeline and isn't guessing at when DTR techniques will appear in consumer drives. ▲

Chill Out With Nanograss

While perhaps not as picturesque as a large field of tall grass swaying gently in the breeze, a breakthrough technology from Bell Labs may soon inspire in you an even more profound sense of wonder. The development is a new method for processing silicon surfaces such that the surfaces look like a field of perfectly level grass blades. Each "nanograss" blade is only a few nanometers in size.

The technology is significant in how it lets a surface interact with liquids. Spilled water, for example, will spatter chaotically on a surface. On nanograss, though, a droplet of liquid rests easily and can be maneuvered at will. Bell research leader Tom Krupenkin coated his nanograss with a nonstick substance, and then he could move droplets without wetting the silicon surface. By adjusting voltages across the silicon surface, Krupenkin's team could sink droplets into the grass and moisten the substrate at will.

"Such behavior may be harnessed to cool computer chips," Krupenkin says in a press statement. "A droplet could be sent to a hot spot on the chip, where it would sink in and absorb the heat, and then go on its way, avoiding the expense and inefficiency of applying a coolant or a heatsink to an entire chip."

Bell is also exploring the potential for using nanograss in optical networking, with the water acting as a switching tool for modifying the properties of light flowing through a data path.

Nanograss might also prove useful in batteries. The silicon could help segregate the liquid electrolyte that's within a battery cell, eliminating any chemical reactions until needed. This would eliminate the usual drain that occurs even when batteries aren't in active use. Because nanograss acts as a water repellant, engineers are talking about using the technology for everything from low-friction torpedoes to self-cleaning windshields. With luck, they might even devise PCs able to withstand coffee and soda pop spills. ▲

by William Van Winkle

Muscle Meets Machine

Like many strange technologies that end up in the most unlikely places, this breakthrough started with a missive from NASA. The organization's Institute for Advanced Concepts wanted a micromachine that could repair tiny meteorite punctures in a spacecraft's surface. The catch was that muscles needed to power the micromachines, not motors. Motors require energy from battery cells—a costly commodity in some environments—while muscles simply need glucose.

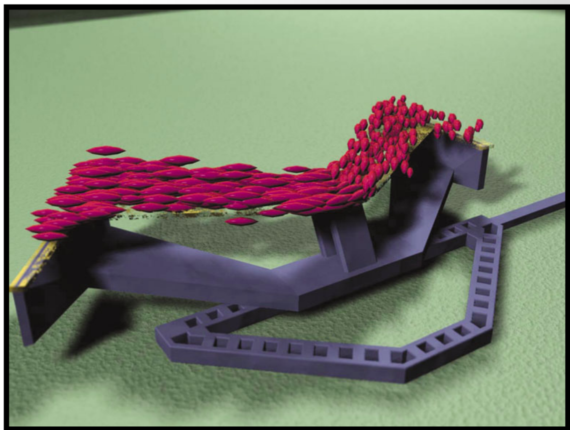

UCLA researcher Carlo Montemagno and his team devised a way for real muscle tissue to replace a motor to create the world's first "musclebots."

Carlos Montemagno and his team at UCLA spent three fruitless years trying countless approaches to accomplish the task. Just recently, though, it got it right. Montemagno and crew carved out a 50-micrometer wide arch of silicon, etched away the coating on the arch's bottom, and replaced it with a gold film that acts as a sort of glue for muscle cells. The arch was left in a Petri dish filled with a glucose culture medium and rat heart cells. Three days later the cells had multiplied and formed muscle fibers joining the gold areas of the arch's underside. When the researchers removed the clamp that held the arch upright during the fiber formation, the "musclebot" started flexing and moving about the dish.

The UCLA researchers admit they are probably decades away from having musclebots swarming over the exterior of a spacecraft to perform repairs. (No one appears to have considered yet how to keep the tissue from freezing while in deep space.) Inter-bot communication also must be considered, but Montemagno, who's also working with Intel to apply his research to enhancing fuel cells, is exploring flexing his musclebots in piezoelectronics. Specifically he wants to craft bots using human muscle tissue with damaged phrenic nerves. This nerve requires very small electric pulses to instruct the diaphragm to help us breath. Musclebots could help supply these signals and alleviate the need for patients to live on ventilators.

"The real near-term application is to be used as a system to investigate muscle atrophy in space," he says. "I am currently working with NASA to incorporate proteomic markers into the muscles and then to evaluate the work that can be done in direct response to protein expression and weightlessness. Using this technology to power implanted devices is also very close to reality. We are currently building the first devices."

Piezoelectronics are also at the heart of some inkjet printing and speaker technologies. Perhaps someday you'll be able to tell friends honestly and literally that your sweet new system packs a lot of muscle. ▲

IBM's New Spin On Semiconductors

While the push in conventional microprocessors is making everything smaller and faster, there's a price to pay in manufacturing expense. But sometimes you don't need speed. It doesn't take a lot of brains to run an RFID tag or a smart card, for example. Yet even the cheapest silicon chips seem expensive when you're talking about adding basic processing functionality into thousands or millions of everyday items. What's needed is a manufacturing process breakthrough that enables basic circuitry on a flexible film and costs only pennies.

Enter IBM's spin coating process, which involves placing several drops of liquid semiconductor material onto a revolving platter. (Visualize putting oil droplets onto a quickly spinning LP.) Centripetal force spreads the liquid into an even coating across the platter's surface. The liquid is then hardened into a thin film that acts as a substrate for transistors and other electronic features. Spin coating isn't actually new, but previous spin coatings have limited the process to only allowing very slow circuits. The advance IBM research leader David Mitzi made led to a 10X increase in circuitry speed that could be obtained in a spin coating process, and further performance gains are still quite possible.

"We're developing ways to process semiconductors over large areas and/or flexible substrates," says Mitzi. "Flexible displays could enable 'electronic paper' applications, and while we haven't studied solar cell applications in particular, the prospects are very exciting, since a key requirement is to lower the cost of the semiconducting cell while maintaining good operating efficiency."

On a chemical level, Mitzi's key innovation was crafting his semiconducting material from hydrazine (a very strong and energetic molecule also used in rocket fuel) with equal amounts of sulfur and tin sulfide. The IBM team's next goal is to find a more benign solvent that still delivers the same results. ▲

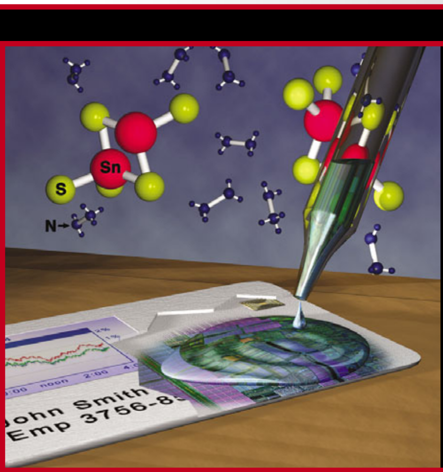

In this conceptual illustration, tin sulfide (red/orange) and hydrazine (blue and green) drip from a pipette, as they do when being applied to a platter for spin coating. The idea is that the end result is extremely cheap electronics for everyday items.

Back Door

Q&A With Jon S. von Tetzchner

Unlike many Web-related CEOs, Jon S. von Tetzchner started at the top and bottom simultaneously. While he co-founded and became the CEO of Opera Software in 1995, he was also one of only two employees—the one who did most of the UI coding. Despite a long and arguably futile battle against Microsoft for the desktop market, von Tetzchner and company continue to build up steam and are jockeying for dominance in the Internet's next frontier: handhelds and set tops.

Q Is most of the excitement over for desktop browsers?

VON TETZCHNER: Well, there is a lot of action. Opera is very good at innovating and finding new features, but when it comes to Web standards, they haven't changed that much because, in a way, IE hasn't changed. IE is setting a lid on development and the way Web pages get sent to users. We're innovating considerably on the mobile and set-top box sides. But most pages are made to be compatible with IE. To a certain extent, there is no way we can show a page better than IE without potentially breaking it.

Q You charge $39 for your browser, although there is a free adware version. What does the $39 buy me?

VON TETZCHNER: Of course, there's no such thing as free lunch, at least not when it comes to IE. You're paying for that in increased prices in Windows. But with us, you're getting more productivity. People have different reasons why they choose Opera: speed of operation, speed of downloading, and the speed of showing pages. They'll talk about multiple windows and keyboard shortcuts. It's just about making things happen quicker and to your liking.

Q What are mouse gestures?

VON TETZCHNER: The idea is that you hold down the right mouse button and then you move the mouse in certain ways to perform actions—going back and forth in history, closing windows, opening windows, etc. You can have these actions do anything that you like.

Q Has the outcome of the Microsoft antitrust trial impacted Opera?

VON TETZCHNER: Well, they were found guilty. But they weren't sent to jail, if you know what I mean. They weren't given a big fine. They weren't told to include Opera or another browser with Windows or make IE optional or any of the natural things they should have been doing. Microsoft changed their licensing agreement so it was theoretically possible for a vendor to include a competing browser. But at the same time, you were already paying for a browser that's included, and the vendors wouldn't risk Microsoft's wrath by doing something against their will.

Q Set tops with built in PVR capabilities are supposed to be huge soon. Will you be competing against TiVo or something else entirely?

VON TETZCHNER: Well, we don't see ourselves as competing with TiVo. I think those kinds of interfaces are great. But we think it makes a lot of sense to use Web technology for all display information, especially information that is being updated frequently, and that applies very much to information in television and related content, things like program guides, something that you'd like to make look a bit better now and then. You'd like to be able to do that without having to program it again, and you can do that using Web technologies.

Q Handheld devies are the next big thing for Opera, so what's next in the handset space?

VON TETZCHNER: Voice has huge potential—being able to get even smaller mobile phones with color screens and still interact with them through voice. And there are a lot of other changes. One thing we've called Opera Platform allows the operator and the user to tailor the phone display with Web content. It's partly coming from the phone—things like your calendar, lost calls, whatever else you might have in the device's applications, and you combine that with our technology. As soon as something interesting happens, based on your interest, you'll get a message on your phone. Not as an SMS or anything like that, but it will appear in your display, potentially with some kind of signal if that's what you prefer.

Readers can see www.cpumag.com/cpujun06/vontetzchner for extra interview content.

William Van Winkle began writing for computer magazines in 1996. He was first published in 1990, the same year he took his first job in computers. He and his family live outside Portland, Ore.

Same award-winning performance.
Now available with a 5-CD Changer.

When we introduced the Acoustic Wave® music system, audio critic Wayne Thompson called it "simply the best-sounding compact stereo system in captivity." *Sound & Vision* reported that it delivers "possibly the best-reproduced sound many people have ever heard." By offering the sound of a big, complicated stereo system – but from an all-in-one music system hardly the size of a briefcase – the Acoustic Wave® music system set a new industry standard.

So how did we make this award-winning system even better? By listening to owners and developing the CD changer they asked for. Engineered specifically for the Acoustic Wave® music system, our 5-CD Changer allows you to enjoy hours of music. You can even use it to connect other audio sources to your system – transforming it into a versatile audio center for movies, music and more.

Lifelike sound – without interruption.

Our separate 5-CD Changer lets you savor the Acoustic Wave® music system's award-winning sound for hours on end. Load your five favorite CDs and press Play. You'll hear your music come to life with crystal-clear highs and deep, enveloping lows that fill the room.

Patented acoustic waveguide speaker technology is what allows the Acoustic Wave® music system to deliver such a spacious, lifelike performance. And our 5-CD Changer lets you enjoy that room-filling sound without ever having to stop the music to change CDs.

Beyond music – without complication.

Our 5-CD Changer also lets you connect your Acoustic Wave® system to additional audio sources and switch among them at the press of a button. Connect it to your TV, VCR, DVD player or computer – and enjoy Bose® quality sound from everything you listen to or watch.

The Acoustic Wave® music system and 5-CD Changer were both designed for complete ease of operation. There are no speakers to connect and no confusing dials to adjust. Every feature operates with intuitive simplicity. And the credit card-sized remote control lets you operate the system and all CD-changer functions with ease.

Try them for yourself, risk free.

Experience the Acoustic Wave® music system and our versatile, easy-to-use 5-CD Changer in your home for 30 days, satisfaction guaranteed. Call today and **save $150** when you order the Acoustic Wave® system together with our 5-CD Changer. Be sure to ask about our **financing options** when you call.

SAVE $150 when you order the Acoustic Wave® music system with our 5-CD Changer by June 30, 2004.

The Acoustic Wave® music system and 5-CD Changer are available from Bose, the most respected name in sound.

Call 1-800-905-1406
Ext. G6252

For information on all our products: **www.bose.com/g6252**

Name____
Address______________________________________
City__________________State______Zip__________
Day Phone________________Eve. Phone______________
E-mail (Optional) _________________________________
Mail to: Bose Corp., Dept. DMG-G6252, The Mountain, Framingham, MA 01701-9168

BOSE®
Better sound through research®

©2004 Bose Corporation. Patent rights issued and/or pending. The Acoustic Wave® music system design is also a registered trademark of Bose Corporation. Financing and reduced price on 5-CD Changer package not to be combined with any other offer or applied to previous purchases, and subject to change without notice. Risk free refers to 30-day trial only. Delivery is subject to product availability. Quotes are reprinted with permission: Wayne Thompson, *Oregonian*, 8/15/89; Julian Hirsch, *Sound & Vision*, 3/85.

Introducing a whole new way to carry your files.

Presenting the go-anywhere, carry-anything ThumbDrive, from Memorex.

Whether you're looking for performance or style, the Memorex ThumbDrive® wins hands down. Incredibly small and durable, this beautifully sculpted drive represents the very best of Flash technology. With capacities ranging from 128MB to 1GB, it transfers tons of photos, music or data at rates up to 50% faster* than competitive USB™ 2.0 products. And the USB connection means it's compatible with both PC and Mac,® as well as many portable electronic devices. The go-anywhere, carry-anything, sleekly-styled ThumbDrive from Memorex—this thumb deserves a hand.

Memorex®

Is it live or is it Memorex?™

THE NO. 1 SELLING BRAND IN DIGITAL MEDIA
No. 1 BRAND

©2003 Memorex Products, Inc. www.memorex.com All trademarks and registered trademarks are of their respective owners

*Based on competitive published information